

Reduced rank regression via adaptive nuclear norm penalization

Kun Chen^{*1}, Hongbo Dong², and Kung-Sik Chan³

¹Department of Statistics, Kansas State University

²Wisconsin Institutes for Discovery, University of Wisconsin Madison

³Department of Statistics and Actuarial Science, University of Iowa

September 25, 2012

Abstract

Adaptive nuclear-norm penalization is proposed for low-rank matrix approximation, by which we develop a new reduced-rank estimation method for the general high-dimensional multivariate regression problems. The adaptive nuclear norm of a matrix is defined as the weighted sum of the singular values of the matrix. For example, the pre-specified weights may be some negative power of the singular values of the data matrix (or its projection in regression setting). The adaptive nuclear norm is generally non-convex under the natural restriction that the weight decreases with the singular value. However, we show that the proposed non-convex penalized regression method has a global optimal solution obtained from an adaptively soft-thresholded singular value decomposition. This new reduced-rank estimator is computationally efficient, has continuous solution path and possesses better bias-variance property than its classical counterpart. The rank consistency and prediction/estimation performance bounds of the proposed estimator are established under high-dimensional asymptotic regime. Simulation studies and an application in genetics demonstrate that the proposed estimator has superior performance to several existing methods. The adaptive nuclear-norm penalization can also serve as a building block to study a broad class of singular value penalties.

1 Introduction

Given n observations of the response $\mathbf{y}_i \in \Re^q$ and predictor $\mathbf{x}_i \in \Re^p$, we consider the multivariate linear regression model:

$$\mathbf{Y} = \mathbf{X}\mathbf{C} + \mathbf{E}, \quad (1.1)$$

*Corresponding author; kunchen@ksu.edu

where $\mathbf{Y} = (\mathbf{y}_1, \dots, \mathbf{y}_n)^T$, $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)^T$, \mathbf{C} is a $p \times q$ coefficient matrix, and $\mathbf{E} = (\mathbf{e}_1, \dots, \mathbf{e}_n)^T$ is a random $n \times q$ matrix of independently and identically distributed random errors with mean zero and variance σ^2 .

We are interested in the scenario when both the number of predictors p and the number of responses q may depend on and even exceed the sample size n . Such high-dimensional regression problems are increasingly encountered in quantitative investigations. It is well known that ordinary least squares (OLS) estimation is equivalent to separately regressing each response on the set of predictors, which, however, ignores the dependence structure of the multivariate response and may be infeasible in high-dimensional settings. The curse of dimensionality may be mitigated by assuming the coefficient matrix \mathbf{C} admit some low-dimensional structure and employing the regularization/penalization approach for model estimation. For example, for Gaussian data, it is appropriate to conduct model estimation by penalized least squares (PLS):

$$\frac{1}{2}\mathcal{J}(\mathbf{C}) + \mathcal{P}_\lambda(\mathbf{C}), \quad (1.2)$$

where $\mathcal{J}(\mathbf{C}) = \|\mathbf{Y} - \mathbf{X}\mathbf{C}\|_F^2$ is the sum of squared error with $\|\cdot\|_F$ denoting the Frobenius norm, $\mathcal{P}_\lambda(\cdot)$ is some penalty function measuring the “size” (complexity) of the enclosed matrix, and λ is a non-negative tuning parameter controlling the degree of penalization.

Within this general framework, an important model is reduced-rank regression (RRR) (Anderson, 1951, 1999, 2002; Izenman, 1975; Reinsel & Velu, 1998), in which dimension reduction is achieved by assuming that the coefficient matrix \mathbf{C} is of low rank, i.e., its rank $r(\mathbf{C}) = r^* < \min(p, q)$. The classical RRR literature mainly focuses on small p cases and maximum likelihood estimation. Bunea et al. (2011) proposed the rank selection criterion (RSC) for high dimensional settings, and revealed that rank constrained estimation can be viewed as a PLS method (1.2) with the penalty proportional to the rank of the coefficient matrix. This l_0 -type penalty can be alternatively cast as a penalty in terms of the number of non-zero singular values of \mathbf{C} , i.e., $\mathcal{P}_\lambda(\mathbf{C}) = \lambda r(\mathbf{C}) = \lambda \sum_k I(d_k(\mathbf{C}) \neq 0)$ where $I(\cdot)$ is the indicator function, which results in an estimator obtained by hard-thresholded singular value decomposition (SVD), see Section 2. Yuan et al. (2007) proposed a nuclear-norm penalized least squares estimator (NNP), in which the nuclear norm penalty is defined as $\mathcal{P}_\lambda(\mathbf{C}) = \lambda \|\mathbf{C}\|_* = \lambda \sum_k d_k(\mathbf{C})$, where $\|\cdot\|_*$ denotes the nuclear norm or the sum of the singular values of the enclosed matrix. This l_1 -type penalty encourages sparsity among the singular values and achieves simultaneous rank reduction and shrinkage coefficient estimation (Negahban & Wainwright, 2011; Bunea et al., 2011; Lu et al., 2012). Rohde & Tsybakov (2011) investigated the theoretical properties of the Schatten- q quasi-norm penalty, which is defined as $\mathcal{P}_\lambda(\mathbf{C}) = \lambda \sum_{k=1}^r d_k^q(\mathbf{C})$ for $0 < q \leq 1$, and nonasymptotic bounds of prediction risk were obtained. Several other methods or theoretical development related to reduced-rank regression exist, see, e.g., Aldrin (2000), Negahban & Wainwright (2011), Mukherjee & Zhu (2011), and Chen et al. (2012). The reduced-rank metrology has connections with many popular tools including principal component analysis and canonical correlation analysis,

and it is extensively studied in matrix completion problems (Candès & Recht, 2009; Candès et al., 2011; Koltchinskii et al., 2011).

It is evident that the aforementioned reduced-rank approaches are closely related to the SVD method in matrix approximation (Eckart & Young, 1936; Reinsel & Velu, 1998). It is also intriguing that the rank and nuclear norm penalized methods can be viewed as l_0 and l_1 penalized methods in the SVD domain, respectively. (In fact, the l_2 norm of the singular values is equivalent to a ridge penalty.) Motivated by these connections, we propose the adaptive nuclear norm regularization method, in which the adaptive nuclear norm (ANN) of a matrix \mathbf{C} is defined as a weighted sum of its singular values; see Xu (2009) for a similar idea related to reweighted penalization in the context of matrix completion. Clearly, the key is to close the gap between the nonsmooth l_0 rank penalty and the l_1 nuclear penalty while keeping the computation stable and efficient for high dimensional data. We show that the convexity of ANN depends on the ordering of the weights. Particularly, the ANN turns out to be non-convex for the case that the weight decreases with the singular value, a condition needed for a meaningful regularization, see Section 2 for further discussion. Despite the non-convexity, we are able to characterize the explicit global optimal solution for the ANN penalized matrix approximation problems.

Based on ANN, we develop a new method of simultaneous dimension reduction and coefficient estimation for the general high-dimensional multivariate regression. Our proposal is based on two main ideas. Firstly, by penalizing the singular values adaptively, our method builds a bridge between the RSC and NNP methods, and it may be viewed as analogous to the adaptive Lasso (Tibshirani, 1996; Zou, 2006) developed for univariate regression. Secondly, the ANN penalty is applied to $\mathbf{X}\mathbf{C}$ rather than \mathbf{C} (Koltchinskii et al., 2011); although the criterion remains non-convex, this setup allows the reduced-rank estimation to be solved explicitly and efficiently. Comparing to the computationally intensive NNP method which tends to overestimate the rank, the proposed ANN method may improve rank determination with the aid of some well-chosen adaptive weights. Comparing to the RSC method, the smooth ANN penalty results in a continuous solution path and allows more flexible bias and variance tradeoff in model fitting. The rank consistency and prediction/estimation performance bounds of the proposed estimator are established under high-dimensional asymptotic regime. We discuss the incorporation of an extra l_2 penalization for improving reduced-rank estimation. Empirical studies demonstrate that the proposed methods enjoy superior performance in both prediction and rank estimation as compared to several existing methods.

2 Adaptive nuclear norm penalty

To study the general properties of the adaptive nuclear norm, we consider the low-rank matrix approximation problem, $\mathbf{Y} = \mathbf{C} + \mathbf{E}$, which is a special case of model (1.1) when \mathbf{X} becomes an identity matrix and $n = p$. In many applications, given the

noisy data matrix \mathbf{Y} , it is of interest to seek its low-rank approximation for denoising, which can be achieved by various methods, e.g., rank penalization or nuclear norm penalization. These methods are closely related to the SVD method in low-rank matrix approximation, which motivated our study.

Consider the SVD of $\mathbf{Y} \in \Re^{n \times q}$,

$$\mathbf{Y} = \mathbf{U}\mathbf{D}\mathbf{V}^T, \quad \mathbf{D} = \text{Diag}\{\mathbf{d}(\mathbf{C})\} = \text{Diag}\{\mathbf{d}\}, \quad (2.1)$$

where \mathbf{U} and \mathbf{V} are respectively $p \times h$ and $q \times h$ orthonormal matrices with $h = \min(p, q)$, and the vector of singular values $\mathbf{d} = (d_1, \dots, d_h)$ consists of non-increasing non-negative singular values of \mathbf{Y} . For any $\lambda \geq 0$, define the hard SVD-thresholding operator (HSVT)

$$\mathcal{H}_\lambda(\mathbf{Y}) = \mathbf{U}\mathcal{H}_\lambda(\mathbf{D})\mathbf{V}^T, \quad \mathcal{H}_\lambda(\mathbf{D}) = \text{Diag}\{d_i I(d_i > \lambda), i = 1, \dots, h\}, \quad (2.2)$$

where $I(\cdot)$ is the indicator function, and the soft SVD-thresholding operator (SSVT)

$$\mathcal{S}_\lambda(\mathbf{Y}) = \mathbf{U}\mathcal{S}_\lambda(\mathbf{D})\mathbf{V}^T, \quad \mathcal{S}_\lambda(\mathbf{D}) = \text{Diag}\{(d_i - \lambda)_+, i = 1, \dots, h\}, \quad (2.3)$$

where x_+ is the non-negative part of x , namely, $x_+ = \max(0, x)$. It is well-known that in the matrix approximation problems ($\mathbf{X} = \mathbf{I}$), an HVT estimator solves (1.2) with the l_0 -type rank penalization (Eckart & Young, 1936), while a SSVT estimator solves (1.2) with the l_1 -type nuclear-norm penalization (Cai et al., 2010); these results are summarized in the proposition below.

Proposition 2.1 *For any $\lambda \geq 0$ and $\mathbf{Y} \in \Re^{n \times q}$, the HVT operator $\mathcal{H}_\lambda(\mathbf{Y})$ defined by (2.2) can be characterized as: $\mathcal{H}_\lambda(\mathbf{Y}) = \arg \min_{\mathbf{C}} \{\|\mathbf{Y} - \mathbf{C}\|_F^2 + \lambda^2 r(\mathbf{C})\}$, and the SSVT operator $\mathcal{S}_\lambda(\mathbf{Y})$ in (2.3) is similarly characterized as: $\mathcal{S}_\lambda(\mathbf{Y}) = \arg \min_{\mathbf{C}} \{\frac{1}{2}\|\mathbf{Y} - \mathbf{C}\|_F^2 + \lambda\|\mathbf{C}\|_*\}$.*

The hard-thresholding operator $\mathcal{H}_\lambda(\mathbf{Y})$ eliminates any singular values below a threshold value λ , while the soft-thresholding operator $\mathcal{S}_\lambda(\mathbf{Y})$ shrinks all the singular values by the same amount λ towards zero. These two SVD operators are natural extensions of the hard/soft-thresholding rules for scalars and vectors (Donoho & Johnstone, 1995; Cai et al., 2010). Generally, estimators based on hard-thresholding often have small bias but may suffer from large variance; in contrast, soft-thresholding reduces variance by introducing extra bias in the estimators, which may be preferable in cases when data are noisy and highly correlated (Donoho & Johnstone, 1995).

The connections between different thresholding rules and penalty terms motivated us to consider the adaptive nuclear norm penalization (ANN). The main idea is to build a bridge between the l_0 and l_1 penalties or the HVT and SSVT rules so as to fine-tune the bias-variance tradeoff in the SVD domain. We define the adaptive nuclear norm of a matrix $\mathbf{C} \in \Re^{p \times q}$:

$$f(\mathbf{C}) = \|\mathbf{C}\|_{*\mathbf{w}} = \sum_{i=1}^h w_i d_i(\mathbf{C}), \quad (2.4)$$

where $h = \min(p, q)$, $d_i(\cdot)$ is the i th largest singular value of the enclosed matrix, and the w_i s are the non-negative weights.

Since the nuclear norm is convex and is a matrix norm, a natural question arises as to whether or not its weighted extension ANN preserves the convexity, which is the case for lasso and adaptive lasso penalties in the vector case (Zou, 2006). However, our analysis shows that the convexity of the ANN depends on the ordering of the non-negative weights. The following theorem gives a necessary and sufficient condition of its convexity.

Theorem 2.2 *For any matrix $\mathbf{C} \in \mathbb{R}^{p \times q}$ ($n = p, h = \min(p, q)$), let $f(\mathbf{C}) = \|\mathbf{C}\|_{*\mathbf{w}}$ defined in (2.4). Then $f(\cdot)$ is convex if and only if $w_1 \geq w_2 \geq \dots \geq w_h \geq 0$.*

Hence, for the ANN to be a convex function, the weight must increase with the singular value, i.e., they are co-monotone. However, to use ANN for penalized estimation, the opposite is desirable, i.e., we would and shall henceforth impose the following order constraint:

$$0 \leq w_1 \leq \dots \leq w_h, \quad (2.5)$$

in order for larger singular values to receive lesser penalty to help reducing the bias and smaller singular values to receive heavier penalty to help promoting sparsity. Here is an example showing that the ANN is neither convex nor concave under constraint (2.5). Consider $n = p = q = 2$, and

$$\mathbf{C}_1 = \begin{pmatrix} 2 & 0 \\ 0 & 1 \end{pmatrix}, \quad \mathbf{C}_2 = \begin{pmatrix} 1 & 0 \\ 0 & 2 \end{pmatrix}.$$

Let $w_1 = 1$ and $w_2 = 2$. It can be verified that $f(\mathbf{C}_1) = f(\mathbf{C}_2) = f(-\mathbf{C}_2) = 4$, while $f((\mathbf{C}_1 + \mathbf{C}_2)/2) = 4.5 > (f(\mathbf{C}_1) + f(\mathbf{C}_2))/2$; also, $f((\mathbf{C}_1 - \mathbf{C}_2)/2) = 1.5 < (f(\mathbf{C}_1) + f(-\mathbf{C}_2))/2$.

The non-convexity of the ANN arises from the constraint (2.5) that the weight decreases with the singular value. In fact, the ANN is then *no longer* a matrix norm. However, we are able to explicitly solve and characterize the global solution of the ANN criterion as follows.

Theorem 2.3 *For any $\lambda \geq 0$, $\mathbf{Y} \in \mathbb{R}^{n \times q}$ and $0 \leq w_1 \leq \dots \leq w_h$ ($n = p, h = \min(n, q)$), a global optimal solution to the optimization problem*

$$\min_{\mathbf{C}} f(\mathbf{C}) := \left\{ \frac{1}{2} \|\mathbf{Y} - \mathbf{C}\|_F^2 + \lambda \sum_{i=1}^h w_i d_i(\mathbf{C}) \right\}. \quad (2.6)$$

is given by the adaptive SVD soft-thresholding (ASVT) operator $\hat{\mathbf{C}} := \mathcal{S}_{\lambda\mathbf{w}}(\mathbf{Y})$,

$$\mathcal{S}_{\lambda\mathbf{w}}(\mathbf{Y}) = \mathbf{U} \mathcal{S}_{\lambda\mathbf{w}}(\mathbf{D}) \mathbf{V}^T, \quad \mathcal{S}_{\lambda\mathbf{w}}(\mathbf{D}) = \text{Diag}\{(d_i - \lambda w_i)_+, i = 1, \dots, h\}, \quad (2.7)$$

Further, if \mathbf{Y} has a unique SVD, $\hat{\mathbf{C}}$ is the unique optimal solution.

The fact that a closed-form global minimizer can be found for the non-convex ANN problem is not immediately clear and rather surprising. The result stems from the von Neumann's trace inequality (Mirsky, 1975) and the properties of SVD, see the Appendix for details. Following Zou (2006), the weights can be set as some power of the singular values of the data matrix, i.e., $\mathbf{w} = \{\mathbf{d}(\mathbf{Y})\}^{-\gamma}$, where $\gamma \geq 0$ is a prespecified constant. In this way, the order constraint (2.5) is automatically satisfied. A more general way of constructing the weights and its relation to other penalty forms will be discussed in Section 7.

3 Adaptive nuclear norm penalization in regression

We now consider the general problem of estimating the coefficient matrix \mathbf{C} , which is possibly of low-rank, in the multivariate linear regression model (1.1). Below, let $\mathbf{P} = \mathbf{X}(\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T$ be the projection matrix onto the column space of \mathbf{X} and $\tilde{\mathbf{C}} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{Y}$ the LS estimator of \mathbf{C} , where $(\mathbf{X}^T \mathbf{X})^{-1}$ denotes the Moore-Penrose inverse of the enclosed Gram matrix. Unless otherwise noted, the singular values are always placed in non-increasing order.

3.1 Rank and nuclear norm penalized regression methods

The fundamental results in Theorem 2.3 about rank and nuclear norm penalization for matrix approximation can be readily extended to the general regression setting. Consider first the rank penalized least squares criterion (Bunea et al., 2011),

$$\frac{1}{2} \|\mathbf{Y} - \mathbf{X}\mathbf{C}\|_F^2 + \lambda r(\mathbf{C}). \quad (3.1)$$

Based on Proposition 2.1, it can be easily shown that the minimizer of (3.1), denoted as $\tilde{\mathbf{C}}^{(\lambda)}$, can be obtained by hard-thresholding the SVD of $\mathbf{X}\tilde{\mathbf{C}}$. Let $\tilde{\mathbf{V}}\tilde{\mathbf{D}}^2\tilde{\mathbf{V}}^T$ be the eigenvalue decomposition of $\mathbf{Y}^T \mathbf{P} \mathbf{Y} = (\mathbf{X}\tilde{\mathbf{C}})^T \mathbf{X}\tilde{\mathbf{C}}$. The SVD of $\mathbf{X}\tilde{\mathbf{C}}$ is then given by $\tilde{\mathbf{U}}\tilde{\mathbf{D}}\tilde{\mathbf{V}}^T$, where $\tilde{\mathbf{U}} = \mathbf{P}\mathbf{Y}\tilde{\mathbf{V}}\tilde{\mathbf{D}}^{-1} = \mathbf{X}\tilde{\mathbf{C}}\tilde{\mathbf{V}}\tilde{\mathbf{D}}^{-1}$. Therefore,

$$\mathbf{X}\tilde{\mathbf{C}}^{(\lambda)} = \mathcal{H}_{\sqrt{2\lambda}}(\mathbf{X}\tilde{\mathbf{C}}) = \mathbf{X}\tilde{\mathbf{C}}\tilde{\mathbf{V}}\tilde{\mathbf{D}}^{-1}\mathcal{H}_{\sqrt{2\lambda}}(\tilde{\mathbf{D}})\tilde{\mathbf{V}}^T, \quad \tilde{\mathbf{C}}^{(\lambda)} = \tilde{\mathbf{C}}\tilde{\mathbf{V}}\tilde{\mathbf{D}}^{-1}\mathcal{H}_{\sqrt{2\lambda}}(\tilde{\mathbf{D}})\tilde{\mathbf{V}}^T. \quad (3.2)$$

This rank selection criterion (RSC) proposed by Bunea et al. (2011) is valid in high-dimensional settings and hence extends the classical rank-constrained RRR approach (Reinsel & Velu, 1998). In fact, the set of rank-constrained estimators which minimize $\|\mathbf{Y} - \mathbf{X}\mathbf{C}\|_F^2$ subject to $r(\mathbf{C}) = r$ ($r = 1, \dots, \min(p, q)$), spans the solution path of (3.1).

The nuclear-norm penalized least squares criterion (NNP) (Yuan et al., 2007)

$$\frac{1}{2} \|\mathbf{Y} - \mathbf{X}\mathbf{C}\|_F^2 + \lambda \|\mathbf{C}\|_*, \quad (3.3)$$

does not have an explicit solution in general, and can be computationally intensive for large-scale data. Extensive research has been devoted to its optimization problem, e.g., Cai et al. (2010), Toh & Yun (2010), etc. One popular algorithm is to iteratively conducting a majorization step of the objective function and a minimization step using soft SVD-thresholding Cai et al. (2010).

The performance of the RSC and NNP estimators is related to the bias-variance trade-off phenomenon discussed in Section 2. The NNP may be more accurate than the RSC when the correlation among predictors is high or the signal to noise ratio (SNR) is low, while RSC may perform better when the correlation is moderate and the SNR is not too low; see Section 6 for details. A drawback of NNP is that it is computationally intensive and is generally not as parsimonious as RSC in rank determination. These motivated our study of the ANN for building a continuum of estimators between the RSC and NNP estimators.

3.2 Adaptive nuclear norm penalized regression method

Predictive accuracy and computation efficiency are both pivotal in high dimensional regression problems. Motivated by criteria (3.1) and (3.3) and their connections with SVD, we propose to estimate \mathbf{C} by minimizing the non-convex PLS criterion

$$\frac{1}{2}\|\mathbf{Y} - \mathbf{X}\mathbf{C}\|_F^2 + \lambda \sum_{i=1}^h w_i d_i(\mathbf{X}\mathbf{C}), \quad (3.4)$$

where $h = \min(p, q)$ and the weights $\{w_i\}$ are required to be non-negative and in non-decreasing order. In practice, a foremost task of using ANN is setting proper adaptive weights. Following Zou (2006), a natural way to construct the weights is based on the LS solution:

$$\mathbf{w} = \{\mathbf{d}(\mathbf{PY})\}^{-\gamma} = \tilde{\mathbf{d}}^{-\gamma}, \quad (3.5)$$

where \mathbf{PY} is the projection of \mathbf{Y} onto the column space of \mathbf{X} and γ a non-negative constant.

The proposed criterion (3.4) is built on two main ideas. Firstly, the criterion directly focuses on prediction matrix approximation and encourages sparsity among the singular values of $\mathbf{X}\mathbf{C}$ rather than those of \mathbf{C} , which may yield low-rank solutions for $\mathbf{X}\mathbf{C}$ and hence for \mathbf{C} . A prominent advantage of this setup is that the problem can then be solved explicitly and efficiently. Secondly, the adaptive penalization of the singular values allows flexible bias-variance tradeoff: large singular values receive small amount of penalization to control the possible bias, and small singular values receive large amount of penalization to induce sparsity and hence reduce the rank. The following Corollary shows that this criterion leads to an explicit ANN estimator.

Corollary 3.1 *A minimizer of (3.4), denoted as $\hat{\mathbf{C}}^{(\lambda\mathbf{w})}$, is obtained via adaptively soft-thresholding the SVD of $\mathbf{X}\tilde{\mathbf{C}}$ where $\tilde{\mathbf{C}}$ is the LS estimator of \mathbf{C} , i.e.,*

$$\mathbf{X}\hat{\mathbf{C}}^{(\lambda\mathbf{w})} = \mathcal{S}_{\lambda\mathbf{w}}(\mathbf{X}\tilde{\mathbf{C}}) = \tilde{\mathbf{U}}\mathcal{S}_{\lambda\mathbf{w}}(\tilde{\mathbf{D}})\tilde{\mathbf{V}}^T, \quad \hat{\mathbf{C}}^{(\lambda\mathbf{w})} = \tilde{\mathbf{C}}\tilde{\mathbf{V}}\tilde{\mathbf{D}}^{-1}\mathcal{S}_{\lambda\mathbf{w}}(\tilde{\mathbf{D}})\tilde{\mathbf{V}}^T. \quad (3.6)$$

where $\tilde{\mathbf{U}}\tilde{\mathbf{D}}\tilde{\mathbf{V}}^T$ is the SVD of $\mathbf{X}\tilde{\mathbf{C}}$ as defined in the previous section.

By Pythagoras' theorem, minimizing the criterion (3.4) is equivalent to minimizing $\{1/2\|\mathbf{X}\tilde{\mathbf{C}} - \mathbf{X}\mathbf{C}\|_F^2 + \lambda \sum_i w_i d_i(\mathbf{X}\mathbf{C})\}$ with respect to \mathbf{C} , where $\tilde{\mathbf{C}}$ is the OLS estimator. The above result then directly follows from Theorem 2.3. The proposed method first projects Y onto the column space of X , i.e., $PY = X\tilde{\mathbf{C}}$, and the ANN estimator is then obtained as a low-rank approximation of PY via soft SVD-thresholding; the thresholding level is adaptive and can be data-driven: the smaller a singular value, the larger its thresholding level. Therefore, the estimated rank of an ANN estimator corresponds to the smallest singular value of PY that exceeds its thresholding level, i.e., $\hat{r} = \max\{r : d_r(PY) > \lambda w_r\}$. For the choice of the weights (3.5), i.e., the estimated rank is given by

$$\hat{r} = \max\{r : d_r(\mathbf{PY}) > \lambda^{\frac{1}{\gamma+1}}\}, \quad (3.7)$$

and the plausible range of the tuning parameter is $\lambda \in [0, \tilde{d}_1^{\gamma+1}]$, with $\lambda = 0$ corresponding to the LS solution and $\lambda = \tilde{d}_1^{\gamma+1}$ the null solution.

The ANN estimator and the RSC estimator only differ in their singular values but the difference can be consequential. While the solution path of RSC is discontinuous and the number of possible solutions equals to the maximum rank, the ANN criterion offers more flexibility in that the resulting solution path is continuous and guided by the data-driven weights. The ANN and RSC are based on the same one-time SVD operation and thus they have similar computation complexity and can both be easily implemented and efficiently computed, in contrast to the computationally intensive NNP method.

For any fixed $\lambda > 0$, the ANN estimator $\hat{\mathbf{C}}^{(\lambda\mathbf{w})}$ can be computed by (3.6). (The same SVD operation can be used to compute the RSC solutions.) To choose an optimal λ and hence an optimal ANN solution, we use the K -fold cross validation (CV) method, based on the predictive performance of the models (Stone, 1974). For the numerical studies reported below, we first compute the solutions over a grid of 100 λ values equally spaced on the log scale and select the best λ value by CV; subsequently we refine the selection process around the chosen λ value with another finer grid of 100 λ values.

4 Rank consistency and error bounds

We study the rank estimation and prediction properties of the proposed ANN estimator. Our theoretical analysis is built on the framework developed by Bunea et al. (2011), as RSC and ANN are closely connected. We mainly focus on the random weights constructed in (3.5), in line with the adaptive Lasso method (Zou, 2006) developed for the univariate (multiple) regression. Similar results are obtained for any prespecified sequence of weights satisfying certain order restriction and boundedness requirements. All the proofs are given in the Appendix.

The rank of the coefficient matrix \mathbf{C} can be viewed as the number of effective combination of predictors linked to the responses. Rank determination is always a foremost task of reduced-rank estimation. The quality of rank estimator, defined by (3.7), clearly depends on the signal to noise ratio. Following Bunea et al. (2011), we shall use the smallest non-zero singular value of \mathbf{XC} , i.e., $d_{r^*}(\mathbf{XC})$, to measure the signal strength, and use the largest singular value of the projected noise matrix \mathbf{PE} , i.e., $d_1(\mathbf{PE})$, to measure that of the noise. Intuitively, if $d_1(\mathbf{PE})$ is much larger than the size of the signal, some signal could be deeply masked by the noise and lost during the thresholding procedure; as such, \hat{r} may be much smaller than the true rank. The lemma below characterizes the “limit” or the true target of \hat{r} and its relationship with the noise level.

Lemma 4.1 *Suppose that there exists an index $s \leq r^*$ such that $d_s(\mathbf{XC}) > (1 + \delta)\lambda^{1/(\gamma+1)}$ and $d_{s+1}(\mathbf{XC}) \leq (1 - \delta)\lambda^{1/(\gamma+1)}$ for some $\delta \in (0, 1]$. Then $P(\hat{r} = s) \geq 1 - P(d_1(\mathbf{PE}) \geq \delta\lambda^{1/(\gamma+1)})$, where \mathbf{P} is the projection matrix onto the column space of \mathbf{X} , \mathbf{E} is the error matrix in model (1.1), and γ is the power parameter in the adaptive weights (3.5).*

This result establishes the relationship between the estimated rank, the signal level, the noise level and the adaptive weights. To achieve consistent rank estimation, we consider the following assumptions:

Assumption The error matrix \mathbf{E} has independent $N(0, \sigma^2)$ entries.

Assumption For any $\theta > 0$, assume $\lambda = \{(1 + \theta)\sigma(\sqrt{r_x} + \sqrt{q})/\delta\}^{\gamma+1}$ with δ defined in Lemma 4.1, and assume $d_{r^*}(\mathbf{XC}) > 2\lambda^{1/(\gamma+1)}$.

Assumption 1 is about the error structure, which ensures that the noise level $d_1(\mathbf{PE})$ is of order $\sqrt{r_x} + \sqrt{q}$, see Lemma .1 (Bunea et al., 2011). Assumption 2 concerns the signal strength relative to the noise level and the appropriate rate of the tuning parameter.

Theorem 4.2 *Suppose Assumptions 1–2 hold. Let $r^* = r(\mathbf{C})$ be the true rank, $r_x = r(\mathbf{X})$ be the rank of \mathbf{X} , and \hat{r} be the estimated rank defined by (3.7). Then $P(\hat{r} = r^*) \rightarrow 1$ as $r_x + q \rightarrow \infty$.*

Theorem 4.2 shows that the ANN estimator is able to identify the correct rank with probability tending to 1 as $r_x + q$ goes to infinity. Similar to Bunea et al. (2011), the consistency results can be extended to the case of sub-Gaussian errors and can also be easily adapted to the case when $r_x + q$ is bounded and the sample size n goes to infinity. Therefore, the rank consistency of the proposed ANN estimator is valid for both classical and high-dimensional asymptotic regimes.

Our main results about the prediction performance of the proposed ANN estimator are presented in Theorem 4.3 below. For simplify, we write $\hat{\mathbf{C}}$ for $\hat{\mathbf{C}}^{(\lambda\mathbf{w})}$.

Theorem 4.3 Suppose Assumptions 1–2 hold. Let $c = d_1(\mathbf{X}\mathbf{C})/d_{r^*}(\mathbf{X}\mathbf{C}) \geq 1$. Then

$$\|\mathbf{X}\hat{\mathbf{C}} - \mathbf{X}\mathbf{C}\|_F^2 \leq \frac{1+a}{1-a} \|\mathbf{X}\mathbf{B} - \mathbf{X}\mathbf{C}\|_F^2 + \frac{1}{a(1-a)} \left\{ \sqrt{2}\delta + 2(2-\delta)^{-\gamma} - (2c+\delta)^{-\gamma} \right\}^2 \lambda^{\frac{2}{\gamma+1}} r^*,$$

with probability greater than $1 - \exp(-\theta^2(r_x+q)/2)$, for any $0 < a < 1$ and any $p \times q$ matrix \mathbf{B} with $r(\mathbf{B}) \leq r^*$. Moreover, taking $\mathbf{B} = \mathbf{C}$ and $a = 1/2$ yields

$$\begin{aligned} \|\mathbf{X}\hat{\mathbf{C}} - \mathbf{X}\mathbf{C}\|_F^2 &\leq 4\{\sqrt{2}\delta + 2(2-\delta)^{-\gamma} - (2c+\delta)^{-\gamma}\}^2 \lambda^{\frac{2}{\gamma+1}} r^* \\ &= 4\{\sqrt{2} + 2(2-\delta)^{-\gamma}/\delta - (2c+\delta)^{-\gamma}/\delta\}^2 (1+\theta)^2 \sigma^2 (\sqrt{r_x} + \sqrt{q})^2 r^* \end{aligned}$$

with probability greater than $1 - \exp(-\theta^2(r_x+q)/2)$.

The above established bound shows that the prediction error is bounded by $d_1^2(\mathbf{PE})r^*$ up to some constant with probability $1 - \exp(-\theta^2(r_x+q)/2)$, i.e., the smaller the error size or the true rank, the smaller the prediction error. The bound is valid for any \mathbf{X} and \mathbf{C} . The estimation error bound of $\hat{\mathbf{C}}$ can also be readily derived from Theorem 4.3, e.g., if $d_{r_x}(\mathbf{X}) \geq \rho > 0$ for some constant ρ , then under Assumptions 1–2, $\|\hat{\mathbf{C}} - \mathbf{C}\|_F^2 \leq 4\rho^{-2}\{\sqrt{2}\delta + 2(2-\delta)^{-\gamma} - (2c+\delta)^{-\gamma}\}^2 \lambda^{\frac{2}{\gamma+1}} r^*$.

The rank consistency and prediction bound can be similarly established for any prespecified sequence of weights satisfying

$$0 \leq w_1 \leq \dots \leq w_{\bar{r}}, w_{r^*} \leq M \leq w_{r^*+1}, \quad (4.1)$$

where $0 < M < \infty$, $w_{r^*+1} > 0$ and $\bar{r} = \min(r_x, q)$; the index s in Lemma 1, the requirements on the tuning sequence and the signal level in Assumption 2 shall be modified accordingly,

$$d_s(\mathbf{X}\mathbf{C}) > (1+\delta)\lambda w_s, d_{s+1}(\mathbf{X}\mathbf{C}) < (1-\delta)\lambda w_{s+1} \text{ for some } \delta \in (0, 1], \quad (4.2)$$

$$\lambda = (1+\theta)\sigma(\sqrt{r_x} + \sqrt{q})(\delta M)^{-1}, d_{r^*}(\mathbf{X}\mathbf{C}) > 2\lambda M. \quad (4.3)$$

Corollary 4.4 Suppose that Assumption 1 and (4.1)–(4.3) hold. Then

- (1) $P(\hat{r} = r^*) \rightarrow 1$ as $r_x + q \rightarrow \infty$;
- (2) $\|\mathbf{X}\hat{\mathbf{C}} - \mathbf{X}\mathbf{C}\|_F^2 \leq \frac{1+a}{1-a} \|\mathbf{X}\mathbf{B} - \mathbf{X}\mathbf{C}\|_F^2 + \frac{1}{a(1-a)} \{(2 + \sqrt{2}\delta)M - w_1\}^2 \lambda^2 r^*$
for any $0 < a < 1$ and B with $r(B) \leq r^*$;
- (3) $\|\mathbf{X}\hat{\mathbf{C}} - \mathbf{X}\mathbf{C}\|_F^2 \leq 4(\sqrt{2} + 2/\delta - w_1/(M\delta))^2 (1+\theta)^2 \sigma^2 (\sqrt{r_x} + \sqrt{q})^2 r^*$
with probability greater than $1 - \exp(-\theta^2(r_x+q)/2)$.

The proof is similar to that of Theorems 4.2 and 4.3 and hence is omitted.

The error bounds of the ANN estimator established in Theorem 4.3 and Corollary 4.4 are comparable to those of the RSC and NNP estimators (Bunea et al., 2011; Rohde & Tsybakov, 2011). The rate of convergence is $(r_x + q)r^*$, which is the optimal minimax rate for rank sparsity under suitable regularity conditions (Rohde & Tsybakov, 2011; Bunea et al., 2012). However, the bounds for NNP was obtained with some extra restrictions on the design matrix, and its tuning sequence that achieves the smallest mean squared error (MSE) usually does not lead to correct rank recovery (Bunea et al., 2011). While both RSC and ANN are able to achieve correct rank recovery and minimal MSE simultaneously, the latter possess continuous solution path with data-driven tuning which may lead to improved empirical performance.

5 Robustification of the reduced-rank estimation

As suggested by a referee and motivated by Mukherjee & Zhu (2011), we discuss the robustification of the reduced-rank methods by incorporating extra l_2 penalty in the penalized criteria.

Mukherjee & Zhu (2011) proposed the robust reduced-rank ridge (RoRR) method which performs l_2 penalized ridge regression under rank constraint. The shrinkage estimation induced by the l_2 penalty makes the reduced rank estimation robust and especially suitable when the predictors are highly correlated. The method can be viewed as the following PLS criterion that was also mentioned in Bunea et al. (2011),

$$\frac{1}{2}\|Y - XC\|_F^2 + \lambda_1 r(C) + \frac{1}{2}\lambda_2 \sum_{i=1}^h d_i^2(C), \quad (5.1)$$

where $h = \min(p, q)$, $\sum_i d_i^2(C) = \text{tr}(C^\top C)$, and λ_1 and λ_2 are tuning parameters. The problem can be solved via data augmentation. Specifically, let

$$Y^* = \begin{pmatrix} Y \\ 0_{p \times q} \end{pmatrix}, \quad X^* = \begin{pmatrix} X \\ \sqrt{\lambda_2} I_{p \times p} \end{pmatrix},$$

then (5.1) can be written as an RSC criterion $1/2\|Y^* - X^*C\|_F^2 + \lambda_1 r(C)$, whose solution is given by (3.2); see Mukherjee & Zhu (2011).

Similarly, the proposed ANN method can also be robustified by incorporating a ridge penalty term. Similar to (3.4), for efficient computation, we impose an l_2 penalty on XC rather than C ,

$$\frac{1}{2}\|Y - XC\|_F^2 + \lambda_1 \sum_{i=1}^h w_i d_i(XC) + \frac{1}{2}\lambda_2 \sum_{i=1}^h d_i^2(XC). \quad (5.2)$$

Interestingly, this robustified ANN criterion (RoANN) is analogous to the adaptive elastic net criterion (Zou & Hastie, 2005) in univariate regression. (The case of imposing

one or both penalties on C directly is more complex and will be pursued elsewhere.) It can be easily verified that, for fixed tuning parameters, the objective function (5.2) is minimized at

$$\hat{C} = \frac{1}{1 + \lambda_2} \hat{C}^{(\lambda_1 w)}$$

where $\hat{C}^{(\lambda_1 w)}$ denotes the ANN estimator in the absence of the l_2 penalty. Indeed, the extra l_2 penalty induces overall shrinkage of the ANN estimator.

For each fixed λ_2 , the RoRR method requires inverting a $p \times p$ matrix of the form $(\mathbf{X}^T \mathbf{X} + \lambda_2 \mathbf{I})$ and performing an SVD of a $q \times q$ matrix. When p is much bigger than n , the Woodbury matrix identity is useful in speeding up computation (Hager, 1989), i.e., $(\mathbf{X}^T \mathbf{X} + \lambda_2 \mathbf{I})^{-1} = 1/\lambda_2 \mathbf{I} - 1/\lambda_2^2 \mathbf{X}^T (\mathbf{I} + 1/\lambda_2 \mathbf{X} \mathbf{X}^T)^{-1} \mathbf{X}$. On the other hand, the RoANN method only requires one-time matrix inversion and SVD operation for obtaining the whole solution path, thereby saving computation. We shall examine the effects of an additional l_2 penalization by simulation in Section 6. Due to space limit, relevant theoretical analysis will be pursued elsewhere.

6 Empirical studies

6.1 Simulation

We compare the prediction, estimation and rank determination performances of the NNP estimator proposed by Yuan et al. (2007), the RSC estimator proposed by Bunea et al. (2011), the RoRR estimator proposed by Mukherjee & Zhu (2011), and our proposed ANN and RoANN estimator. In ANN estimation the adaptive weights are constructed as (3.5) with $\gamma = 0, 1, 2$ and we denote the resulting estimator as ANN_γ ($\gamma = 0$ means unweighted ANN). In the numerical results reported below, we use the accelerated proximal gradient algorithm implemented in Matlab by Toh & Yun (2010) for NNP estimation. R code for RoRR was provided by their original authors (Mukherjee & Zhu, 2011), and we modified their code to make use of the Woodbury matrix identity (Hager, 1989) for saving computation. We have also implemented all the other methods in R (R Development Core Team, 2008). All computation was done on linux machines with 3.4 GHz CPU and 8 GB RAM.

We consider the same simulation models as in Bunea et al. (2011). The coefficient matrix \mathbf{C} is constructed as $\mathbf{C} = b \mathbf{C}_0 \mathbf{C}_1^T$, where $b > 0$, $\mathbf{C}_0 \in \Re^{p \times r^*}$, $\mathbf{C}_1 \in \Re^{q \times r^*}$ and all entries in \mathbf{C}_0 and \mathbf{C}_1 are i.i.d. $N(0, 1)$. Two scenarios of model dimensions are considered, i.e., $p, q < n$ and $p, q > n$.

- Model I ($n = 100$, $p = q = 25$, $r^* = 10$): The covariate matrix \mathbf{X} is constructed by generating its n rows as i.i.d. samples from a multivariate normal distribution $\text{MVN}(\mathbf{0}, \Gamma)$, where $\Gamma = (\Gamma_{ij})_{p \times p}$ and $\Gamma_{ij} = \rho^{|i-j|}$ with some $0 < \rho < 1$.
- Model II ($n = 20$, $p = q = 25$, $r^* = 5$, $r_x = 10$): The covariate matrix \mathbf{X} is generated as $\mathbf{X} = \mathbf{X}_0 \Gamma^{1/2}$, where Γ is defined as above, $\mathbf{X}_0 = \mathbf{X}_1 \mathbf{X}_2$, $\mathbf{X}_1 \in \Re^{n \times r_x}$, $\mathbf{X}_2 \in \Re^{r_x \times p}$, and all entries of \mathbf{X}_1 , \mathbf{X}_2 are i.i.d $N(0, 1)$.

The data matrix \mathbf{Y} is then generated by $\mathbf{Y} = \mathbf{X}\mathbf{C} + \mathbf{E}$, where the elements of \mathbf{E} are i.i.d. samples from $N(0, 1)$. It can be seen that each simulated model is characterized by the following parameters : n (sample size), p (number of predictors), q (number of responses), r^* (rank of \mathbf{C}), r_x (rank of \mathbf{X}), ρ (design correlation), and b (signal strength). The experiment was replicated 500 times for each parameter setting.

To alleviate the influence of the inaccuracy in empirical tuning parameter selection and to reveal the true potential of each penalized method for fair comparison, one way is to tune each method based on its prediction accuracy evaluated on a very large independently generated validation data set; this yields “optimally tuned” estimators denoted as $NNP^{(O)}$, $RSC^{(O)}$, $ANN_{\gamma}^{(O)}$, etc. The 10-fold cross validation method is also used with the actual data, which results in the estimators $NNP^{(C)}$, $RSC^{(C)}$, $ANN^{(C)}$, etc. (Although to save space, we only report the CV results of RSC and ANN_2 .) For each method, the model accuracy is measured by the average of the scaled mean-squared-errors (SMSE) from all runs, i.e., $SMSE = 100\|\mathbf{C} - \hat{\mathbf{C}}\|_F^2/(pq)$ for estimation (Est), and $SMSE = 100\|\mathbf{X}\mathbf{C} - \mathbf{X}\hat{\mathbf{C}}\|_F^2/(nq)$ for prediction (Pred). Their standard errors are also reported. To evaluate the rank determination performance, we report (1) the average of the estimated ranks from all runs, and (2) the percentage of correct rank identification. For each method, the average computation time per replication is reported.

Tables 1 and 2 report the simulation results for Model I and II, respectively. We summarize our findings as follows.

- We first examine the performance of ANN_{γ} with different γ values. The ANN_0 , which does not use the adaptive weights, is not as accurate as the ANN with adaptive weights in both rank estimation and prediction in general. Its behavior is similar to that of the NNP: they both tend to overestimate the rank but may perform well when the signal is weak and the model dimension is not high (Model I). The performance of both ANN_1 and ANN_2 is substantially better than ANN_0 with the aid of adaptive weights. We have also experimented with other γ values, and as expected, the ANN estimator behave more and more similar to the RSC as γ increases. Our results show that $\gamma = 2$ is generally a very good choice. Henceforth we always refer to ANN_2 in the following comparisons with other reduced-rank methods.
- The ANN_2 generally outperforms the RSC in both estimation and prediction. The improvement can be substantial, especially when the signal is weak or moderate and the correlation among the predictors is high. For rank determination, both the ANN and the RSC have similar excellent behavior when the signal is moderate to strong and the correlation among the predictors is weak to moderate. We notice that the ANN estimator tends to slightly overestimate the rank; however, the overestimation is generally negligible. When the signal strength is moderate to large and correlation among predictors is weak to moderate ANN does slightly worse in terms of rank selection than RSC yet is able to maintain a small gain in estimation and prediction. This gain is due to the shrinkage effect

Table 1: Comparison of estimation, prediction and rank determination performances of various reduced-rank estimators using Model I ($n = 100, p = 25, q = 25, r^* = 10$). The superscript (O) stands for optimal tuning, and (C) stands for cross validation. The estimation error (Est) and prediction error (Pred) are reported along with their standard errors in the parenthesis. For rank estimation (Rank), the average of estimated rank and the percentage of correct rank identification are reported. The simulation is based on 500 replications, and the average running time of each replication is reported in seconds (Time).

b	ERR	Method							
		NNP ^I	RSC ^O	RSC ^C	ANN ₂ ^O	ANN ₂ ^C	ANN ₁ ^O	ANN ₀ ^O	RoRR ^O
$\rho = 0.9$									
0.05	Est	1.57 (0.2)	3.20 (0.6)	3.23 (0.7)	2.61 (0.4)	2.67 (0.5)	2.48 (0.4)	2.46 (0.3)	1.60 (0.2)
	Pred	7.82 (0.8)	12.22 (1.3)	12.69 (1.6)	9.94 (1)	10.22 (1.1)	9.68 (1)	10.48 (1)	8.06 (0.8)
	Rank	7.66, 2.8	3.28, 0	3.10, 0	5.53, 0	4.88, 0	6.71, 0	10.89, 25.6	7.57, 3.6
0.1	Est	3.53 (0.4)	5.62 (0.7)	5.72 (0.8)	4.52 (0.5)	4.61 (0.6)	4.29 (0.5)	4.31 (0.5)	3.61 (0.4)
	Pred	12.07 (1)	16.07 (1.4)	16.56 (1.6)	13.54 (1.2)	13.78 (1.2)	13.15 (1.1)	14.36 (1.1)	12.48 (1)
	Rank	11.18, 18.6	6.21, 0	5.95, 0.2	8.12, 4.6	7.64, 2.2	9.23, 33.2	14.06, 0	7.85, 5.2
0.2	Est	5.59 (0.6)	6.84 (0.7)	7.01 (0.8)	5.88 (0.6)	5.97 (0.6)	5.67 (0.6)	6.02 (0.6)	5.54 (0.6)
	Pred	15.5 (1.1)	17.17 (1.3)	17.58 (1.4)	15.53 (1.1)	15.7 (1.2)	15.27 (1.1)	17.23 (1.3)	15.36 (1.1)
	Rank	10.87, 20.8	9.00, 26.8	8.78, 19.8	10.03, 58.8	9.71, 55.2	10.95, 26	16.25, 0	9.38, 47.2
0.3	Est	6.40 (0.7)	6.83 (0.8)	6.94 (0.8)	6.22 (0.7)	6.29 (0.7)	6.11 (0.7)	6.81 (0.7)	6.05 (0.7)
	Pred	16.40 (1.2)	16.86 (1.2)	17.09 (1.4)	15.96 (1.2)	16.06 (1.2)	15.88 (1.2)	18.51 (1.2)	15.91 (1.2)
	Rank	10.69, 30.2	9.81, 81.4	9.67, 68.4	10.33, 64.2	10.14, 73.6	11.2, 10.8	17.00, 0	9.86, 86
$\rho = 0.5$									
0.05	Est	0.80 (0.1)	1.19 (0.1)	1.24 (0.1)	0.92 (0.1)	0.94 (0.1)	0.87 (0.1)	0.85 (0.1)	0.84 (0.1)
	Pred	12.36 (1)	16.84 (1.5)	17.77 (2)	13.56 (1.2)	13.80 (1.3)	12.97 (1.1)	13.25 (1.1)	12.96 (1.1)
	Rank	12.83, 0.8	6.02, 0	5.63, 0	8.01, 5.2	7.60, 1.8	9.18, 30.4	13.28, 0	7.88, 6.4
0.1	Est	1.19 (0.1)	1.35 (0.1)	1.38 (0.2)	1.17 (0.1)	1.19 (0.1)	1.13 (0.1)	1.20 (0.1)	1.15 (0.1)
	Pred	16.21 (1.2)	17.42 (1.2)	17.82 (1.5)	15.72 (1.1)	15.86 (1.2)	15.37 (1.1)	16.82 (1.1)	15.74 (1.1)
	Rank	15.21, 0	9.15, 32	8.93, 24.6	10.02, 57.8	9.80, 56.4	10.98, 23.8	15.83, 0	9.39, 47.2
0.2	Est	1.36 (0.1)	1.29 (0.1)	1.29 (0.1)	1.23 (0.1)	1.24 (0.1)	1.22 (0.1)	1.41 (0.1)	1.23 (0.1)
	Pred	17.63 (1.4)	16.59 (1.2)	16.67 (1.2)	16.08 (1.1)	16.16 (1.1)	16.1 (1.1)	19.00 (1.3)	16.15 (1.1)
	Rank	12.59, 0.4	9.98, 98.2	9.95, 94	10.28, 72.2	10.15, 85	11.16, 10.2	17.34, 0	9.99, 98.8
0.3	Est	1.32 (0.1)	1.25 (0.1)	1.25 (0.1)	1.22 (0.1)	1.23 (0.1)	1.22 (0.1)	1.47 (0.1)	1.22 (0.1)
	Pred	17.08 (1.3)	16.22 (1.2)	16.22 (1.2)	15.95 (1.2)	16.03 (1.2)	16.00 (1.2)	19.47 (1.3)	16.00 (1.2)
	Rank	10.92, 14	10.00, 99.8	10.00, 99.4	10.21, 80.4	10.12, 89.8	11.04, 17	17.83, 0	10.00, 99.8
$\rho = 0.1$									
0.05	Est	0.64 (0.1)	0.86 (0.1)	0.90 (0.1)	0.69 (0.1)	0.70 (0.1)	0.65 (0.1)	0.65 (0.1)	0.65 (0.1)
	Pred	13.25 (1)	17.39 (1.4)	18.2 (1.8)	14.12 (1.2)	14.41 (1.3)	13.50 (1.1)	13.71 (1.1)	14.01 (1.2)
	Rank	14.28, 0	6.64, 0.6	6.21, 0	8.48, 12	7.95, 5.6	9.64, 39.8	13.62, 0	8.10, 10
0.1	Est	0.85 (0.1)	0.90 (0.1)	0.91 (0.1)	0.81 (0.1)	0.82 (0.1)	0.79 (0.1)	0.85 (0.1)	0.81 (0.1)
	Pred	16.82 (1.2)	17.32 (1.3)	17.65 (1.4)	15.88 (1.2)	16.01 (1.2)	15.59 (1.2)	17.16 (1.2)	16.05 (1.2)
	Rank	16.51, 0	9.46, 52	9.27, 40	10.16, 63.6	9.92, 64.8	11.13, 15.4	16.09, 0	9.61, 63.2
0.2	Est	0.97 (0.1)	0.85 (0.1)	0.85 (0.1)	0.82 (0.1)	0.83 (0.1)	0.82 (0.1)	0.96 (0.1)	0.82 (0.1)
	Pred	18.57 (1.4)	16.39 (1.2)	16.44 (1.2)	15.99 (1.1)	16.07 (1.1)	16.00 (1.1)	18.99 (1.3)	16.02 (1.2)
	Rank	15.36, 0	10.00, 100	9.98, 97	10.22, 78.4	10.18, 84.8	11.09, 12	17.38, 0	10.00, 100
0.3	Est	0.92 (0.1)	0.83 (0.1)	0.83 (0.1)	0.82 (0.1)	0.83 (0.1)	0.82 (0.1)	1.00 (0.1)	0.82 (0.1)
	Pred	17.7 (1.5)	16.15 (1.2)	16.17 (1.2)	15.97 (1.2)	16.05 (1.2)	16.03 (1.2)	19.64 (1.3)	15.99 (1.2)
	Rank	11.25, 3	10.00, 100	10.01, 99.4	10.18, 82.8	10.14, 87.6	10.95, 17.2	17.86, 0	10.00, 100
Time		17.52	0.02	0.02	0.15	0.18	0.15	0.15	3.93
									2.06

Table 2: Comparison of estimation, prediction and rank determination performances of various reduced-rank estimators using Model II ($n = 20, p = 100, q = 25, r^* = 5, r_x = 10$). The layout of the table is the same as in Table 1.

b	ERR	NNP ^O	RSC ^O	RSC ^C	ANN ^O ₂	Method ANN ^C ₂	ANN ^O ₁	ANN ^O ₀	RoRR ^O	RoANN ^O ₂
$\rho = 0.9$										
0.05	Est	1.14 (0.2)	1.15 (0.2)	1.15 (0.2)	1.15 (0.2)	1.15 (0.2)	1.15 (0.2)	1.15 (0.2)	1.15 (0.2)	1.15 (0.2)
	Pred	34.12 (4.1)	31.49 (4)	32.32 (4.6)	29.08 (3.5)	29.65 (3.5)	28.95 (3.4)	35.28 (4.2)	29.75 (3.5)	28.89 (3.5)
	Rank	7.38, 1.6	4.73, 73.4	4.63, 64.4	5.19, 71.2	4.99, 73.2	5.70, 36.4	8.26, 0	4.83, 83.2	5.30, 65.8
0.1	Est	4.50 (0.7)	4.52 (0.7)	4.52 (0.7)	4.52 (0.7)	4.52 (0.7)	4.52 (0.7)	4.52 (0.7)	4.52 (0.7)	4.52 (0.7)
	Pred	37.03 (4.4)	30.83 (3.6)	31.00 (3.8)	30.01 (3.6)	30.40 (3.8)	30.15 (3.7)	38.86 (4.8)	30.25 (3.6)	29.88 (3.6)
	Rank	7.57, 2.8	4.99, 99.2	4.97, 96.6	5.22, 79.4	5.13, 87.4	5.63, 42.8	8.77, 0	4.99, 99.4	5.27, 74.6
0.2	Est	18.06 (2.7)	17.92 (2.6)	17.92 (2.6)	17.92 (2.6)	17.92 (2.6)	17.93 (2.6)	17.92 (2.6)	17.92 (2.6)	17.92 (2.6)
	Pred	34.90 (5.5)	29.79 (3.6)	29.81 (3.6)	29.55 (3.6)	29.81 (3.6)	29.67 (3.6)	40.08 (4.8)	29.59 (3.6)	29.52 (3.6)
	Rank	6.33, 39.4	5.00, 100	5.00, 99.8	5.22, 79.2	5.12, 89	5.53, 51.4	9.02, 0	5.00, 100	5.29, 72.8
0.3	Est	40.42 (5.6)	40.58 (5.9)	40.58 (5.9)	40.58 (5.9)	40.58 (5.9)	40.58 (5.9)	40.58 (5.9)	40.58 (5.9)	40.58 (5.9)
	Pred	32.29 (5.5)	30.18 (3.5)	30.22 (3.6)	30.06 (3.5)	30.32 (3.6)	30.16 (3.5)	41.29 (5)	30.04 (3.5)	30.03 (3.5)
	Rank	5.48, 75	5.00, 100	5.00, 99.6	5.18, 82	5.15, 87.4	5.46, 56.2	9.07, 0	5.00, 100	5.26, 75.2
$\rho = 0.5$										
0.05	Est	1.12 (0.2)	1.12 (0.2)	1.12 (0.2)	1.12 (0.2)	1.12 (0.2)	1.12 (0.2)	1.12 (0.2)	1.12 (0.2)	1.12 (0.2)
	Pred	34.98 (4.2)	31.01 (3.9)	31.57 (4.2)	29.14 (3.6)	29.65 (3.8)	29.12 (3.6)	35.92 (4.7)	29.73 (3.6)	29.00 (3.6)
	Rank	7.97, 0	4.90, 89.8	4.79, 78	5.20, 77	5.06, 79.2	5.68, 37.6	8.43, 0	4.93, 93.2	5.32, 67.8
0.1	Est	4.50 (0.6)	4.53 (0.6)	4.53 (0.6)	4.53 (0.6)	4.53 (0.6)	4.53 (0.6)	4.53 (0.6)	4.53 (0.6)	4.53 (0.6)
	Pred	37.83 (4.6)	30.16 (3.5)	30.29 (3.6)	29.58 (3.3)	29.93 (3.6)	29.73 (3.4)	38.98 (4.2)	29.76 (3.4)	29.52 (3.4)
	Rank	7.76, 7.2	4.99, 99.4	4.99, 98	5.19, 81	5.10, 90	5.63, 42.2	8.86, 0	4.99, 99.4	5.27, 73.6
0.2	Est	17.90 (2.6)	18.03 (2.7)	18.03 (2.7)	18.03 (2.7)	18.03 (2.7)	18.03 (2.7)	18.04 (2.7)	18.03 (2.7)	18.03 (2.7)
	Pred	34.84 (6.1)	30.05 (3.6)	30.07 (3.6)	29.89 (3.6)	30.15 (3.6)	29.99 (3.6)	40.75 (5)	29.88 (3.6)	29.85 (3.6)
	Rank	6.25, 51.4	5.00, 100	5.00, 99.8	5.21, 79.8	5.14, 87.8	5.51, 53.4	9.07, 0	5.00, 100	5.27, 75
0.3	Est	40.44 (6.1)	40.13 (6)	40.13 (6)	40.13 (6)	40.13 (6)	40.13 (6)	40.13 (6)	40.13 (6)	40.13 (6)
	Pred	31.55 (4.9)	30.32 (3.1)	30.32 (3.1)	30.23 (3.1)	30.44 (3.1)	30.3 (3.1)	41.30 (4)	30.18 (3.1)	30.2 (3.1)
	Rank	5.34, 85.6	5.00, 100	5.00, 100	5.23, 77.8	5.11, 89	5.46, 56.8	9.16, 0	5.00, 100	5.32, 70.2
$\rho = 0.1$										
0.05	Est	1.13 (0.2)	1.13 (0.2)	1.13 (0.2)	1.13 (0.2)	1.13 (0.2)	1.13 (0.2)	1.13 (0.2)	1.13 (0.2)	1.13 (0.2)
	Pred	35.17 (4.2)	31.02 (3.8)	31.86 (4.5)	29.25 (3.5)	29.7 (3.6)	29.24 (3.5)	36.04 (4.3)	29.85 (3.5)	29.10 (3.6)
	Rank	8.07, 0	4.88, 88.2	4.77, 76.8	5.19, 75.6	5.10, 78.8	5.70, 35	8.41, 0	4.91, 91	5.32, 65.4
0.1	Est	4.55 (0.7)	4.47 (0.7)	4.47 (0.7)	4.47 (0.7)	4.47 (0.7)	4.47 (0.7)	4.47 (0.7)	4.47 (0.7)	4.47 (0.7)
	Pred	37.73 (4.6)	30.24 (3.5)	30.33 (3.7)	29.72 (3.5)	30.04 (3.5)	29.81 (3.5)	38.88 (4.4)	29.8 (3.5)	29.64 (3.4)
	Rank	7.89, 5.4	5.00, 99.8	4.99, 99.2	5.22, 79.4	5.12, 89.2	5.62, 43.4	8.84, 0	5.00, 100	5.29, 72.6
0.2	Est	18.1 (2.6)	17.92 (2.6)	17.92 (2.6)	17.92 (2.6)	17.92 (2.6)	17.92 (2.6)	17.92 (2.6)	17.92 (2.6)	17.92 (2.6)
	Pred	34.48 (5.9)	30.13 (3.4)	30.13 (3.4)	29.98 (3.4)	30.24 (3.4)	30.09 (3.4)	40.76 (4.7)	29.97 (3.4)	29.93 (3.4)
	Rank	6.18, 54.4	5.00, 100	5.00, 100	5.23, 78	5.13, 88.2	5.51, 52.6	9.07, 0	5.00, 100	5.30, 71.6
0.3	Est	40.48 (6.2)	40.48 (5.6)	40.48 (5.6)	40.48 (5.6)	40.48 (5.6)	40.48 (5.6)	40.48 (5.6)	40.48 (5.6)	40.48 (5.6)
	Pred	31.31 (5.1)	30.06 (3.4)	30.11 (3.5)	29.96 (3.4)	30.28 (3.5)	30.05 (3.4)	41.17 (4.6)	29.91 (3.4)	29.94 (3.4)
	Rank	5.31, 86.2	5.00, 100	5.00, 99.6	5.24, 76	5.15, 87	5.46, 55.2	9.08, 0	5.00, 100	5.34, 67.8
Time		20.22	0.02	0.08	0.15	0.26	0.15	0.15	3.95	2.49

using soft-thresholding. While RSC is based on hard-thresholding and keeps the first few leading SVD layers completely as signal, in reality all the SVD layers are contaminated by noise, and hence the shrinkage estimation is effective. When the signal is weak and the correlation among the predictors is high, the rank determination performance of the RSC can be much worse than that of the ANN, which is also reflected in their differences in prediction and estimation.

- The ANN_2 also outperforms the NNP in general, and is more parsimonious than the NNP in both rank reduction and computation. Only when the signal is weak and the correlation among predictors is very high, may the NNP method slightly outperform the ANN (and hence RSC) in estimation and prediction. However, this gain exacts a high cost that may make it not worthwhile, as the NNP often excessively overestimates the rank and is much harder to compute. Our findings regarding the NNP agrees with those reported in [Bunea et al. \(2011\)](#). Note that [Bunea et al. \(2011\)](#) also proposed a calibrated NNP method for rank determination, which showed very similar behavior as that of the RSC. (Hence it is not reported here.)
- Adding an l_2 penalty usually boosts the performance of the reduced-rank methods. The RoRR may substantially outperform its non-robustified counterpart RSC, especially when the correlation is high, showing the power of shrinkage estimation. The RoANN₂ only slightly outperforms ANN₂, because ANN₂ itself performs adaptive shrinkage estimation and thus the gain from the overall shrinkage induced by the l_2 penalty is much limited. RoRR, ANN₂ and RoANN₂ have comparable performance in most cases and are the best methods.
- NNP is much more computationally expensive than the other reduced-rank methods. Both RSC and ANN are very fast to compute, and in practice they can always be computed together as they rely on the same SVD operation. Adding an l_2 penalty increases computation time. RoRR takes longer computation time than RoANN, for the reason explained in Section 5.

Overall, the ANN approach is preferable to both the RSC and the NNP methods, especially for cases when the data are noisy and the correlation among the predictors is high. Adding extra l_2 penalty is certainly worthwhile, especially for RSC, but it may incur a lot of computation efforts.

6.2 An application in Genomics

We consider a breast cancer data set ([Witten et al., 2009; Bunea et al., 2010](#)), which consists of the gene expression measurements and the comparative genomic hybridization (CGH) measurements for $n = 89$ subjects. The data set is available in the R package *PMA*, and a detailed description can be found in [Chin et al. \(2006\)](#).

Table 3: Performances of reduced-rank estimators on breast cancer data. Setting 1 regresses GEPs on CNVs, and setting 2 regresses CNVs on GEPs.

	Method				
	NNP ^C	RSC ^C	RoR ^C	ANN ₂ ^C	RoANN ₂ ^C
Setting 1: Chromosome 14 ($q = 641, p = 76$)					
CVE	0.84	1.00	0.85	0.96	0.92
Rank	4	0	7	1	1
Time	2190.4	5.5	689.0	24.6	38.2
Setting 2: Chromosome 14 ($q = 76, p = 641$)					
CVE	0.69	0.59	0.59	0.58	0.58
Rank	23	5	5	11	17
Time	2158.5	13.5	186.7	13.8	16.8
Setting 1: Chromosome 21 ($q = 227, p = 44$)					
CVE	0.84	0.95	0.83	0.87	0.86
Rank	3	1	2	1	1
Time	1266.4	0.5	66.2	2.6	8.4
Setting 2: Chromosome 21 ($q = 44, p = 227$)					
CVE	0.69	0.65	0.63	0.62	0.62
Rank	5	1	1	1	2
Time	1226.0	0.5	28.6	0.7	2.9

Prior studies have demonstrated a link between DNA copy-number changes and cancer risk (Pollack et al., 2002; Peng et al., 2010). It is thus of interest to examine the relationship between DNA copy number variations (CNVs) and gene expression profiles (GEPs), for which multivariate regression methods can be useful. Biologically, it makes sense to regress GEPs on CNVs since the latter play an important role in regulating the former. The reverse approach of regressing CNVs on GEPs is also meaningful, in that the resulting predictive model may identify functionally relevant CNAs; this approach has been shown to be promising in enhancing the limited CGH analysis with the wealth of GEP data (Geng et al., 2011; Zhou et al., 2012). We thus try both approaches, i.e., setting 1: designate the CNVs on the CGH spots of a chromosome as predictors ($\mathbf{X}_{n \times p}$), and the GEPs of the same chromosome as responses ($\mathbf{Y}_{n \times q}$), and setting 2: reverse the roles of \mathbf{X} and \mathbf{Y} . Both the responses and predictors are centered and standardized.

We focus on chromosomes 14 and 21, which were previously studied in Bunea et al. (2010). In this study, p and q are either comparable or much larger than $n = 89$. To alleviate the high-dimensionality problem, the reduced-rank methods are appealing for identifying a few linear combinations of predictors for optimally predicting the response variables. We perform model estimation using various reduced-rank methods, with the tuning parameters selected by ten-fold cross validation. (We obtained similar results with five-fold and fifteen-fold cross validations.) The cross validation error rate (CVE) (averaged over the number of response variables and the sample size) is then used to compare the predictive performance of different penalization schemes. Note that CVE = 1 corresponds to a null model with zero coefficient matrix. Table 3 reports the CVE, the estimated rank and the computation time.

In setting 1, the SNR is very low as reflected by the CVEs being close to 1. The RSC may fail to pick up any signal (for chromosome 14), while the other methods, especially NNP and RoRR, perform better owing to the power of shrinkage estimation. In setting 2, the SNR is relatively higher, and the ANN and RoANN have better

prediction performance than all the other methods. In Section 4, we have shown that the prediction error is of the order $(r_x + q)r^*$, where $r_x \leq \min(n, p)$ is the rank of the design matrix \mathbf{X} . This may partly explain why in setting 1 the prediction is always poorer than in setting 2, because q is much larger than p when the CNVs (GEPs) serve as predictors (responses). For rank estimation, the NNP method always yields higher rank estimates than the other methods; the ANN estimated rank is higher than that of RSC for chromosome 14 in setting 2, otherwise their rank estimates are similar. Adding l_2 penalty improves the predictive performance of the reduced-rank methods, especially the improvement from RSC to RoRR can be substantial. The robustified methods may also yield higher rank estimates. As borne out by the simulation study, such behaviors result from the hybridization between reduced-rank methods and ridge regression. Both NNP and RoRR are computationally intensive for large datasets, while RSC, ANN and RoANN are much faster to compute. Overall, it can be seen that the proposed ANN approach shows better performance than the RSC, with not much extra computational cost.

7 Discussion

There are several potential directions for future research. We have mainly considered the adaptive nuclear-norm penalization on \mathbf{XC} which yields a computationally efficient SVD-thresholding estimator for dimension reduction and shrinkage estimation. It is interesting to consider an ANN criterion that puts the adaptive nuclear-norm penalization directly on the coefficient matrix \mathbf{C} , i.e., $\{\|\mathbf{Y} - \mathbf{XC}\|_F^2 + \lambda \sum w_i d_i(\mathbf{C})\}$. This criterion may be optimized by similar iterative SVD-thresholding method used in solving the NNP problem. Although the computation will be more intensive, the advantage is that this criterion would result in simultaneous adaptive rank reduction and shrinkage estimation.

The proposed ANN method can serve as the building block to study a family of singular value penalties. This is based on the connection between an adaptive l_1 -type penalty and many concave penalty functions such as SCAD (Fan & Li, 2001) and bridge penalty (Knight & Fu, 2000). Consider the general regression problem (1.2) with a general singular value penalty $\mathcal{P}_\lambda(\mathbf{C}) = \sum_{i=1}^h p_\lambda(d_i)$, where $p_\lambda(\cdot)$ is a penalty function, e.g., l_q bridge penalty $p_\lambda(|d_i|) = \lambda |d_i|^q$ ($0 < q < 1$) (Rohde & Tsybakov, 2011). In this setup the optimization of (1.2) can be challenging. A promising approach is to adopt a local linear approximation (Zou & Li, 2008), $p_\lambda(|d_i|) \approx p_\lambda(|d_i^{(0)}|) + p'_\lambda(|d_i^{(0)}|)(|d_i| - |d_i^{(0)}|)$, for $d_i \approx d_i^{(0)}$, where $d_i^{(0)}$ is some initial estimator of d_i which for example can be obtained by the LS method. It can be seen that for fixed $d_i^{(0)}$, up to a constant, the first-order approximated penalty admits exactly an ANN form. This suggests the ANN estimator with the weights $p'_\lambda(|d_i^{(0)}|)$ can be viewed as an one-step estimator of these problems, and these problems may be solved by an iteratively reweighted ANN approach.

It is shown that incorporating an extra ridge penalty can induce further shrinkage and hence improve the reduced-rank estimation. When combined with the ANN

penalty, such a criterion bears resemblance to the elastic-net criterion (Enet) (Zou & Hastie, 2005) in univariate regression. It would be interesting to investigate the properties of the SVD-Enet approaches. Another pressing problem concerns further extending the regularized reduced-rank regression methods to generalized linear models and nonparametric regression models (Yee & Hastie, 2003; Li & Chan, 2007). On the optimization aspect, it is interesting to study the usage of ANN in some classical sparse optimization areas, such as matrix completion (Candès et al., 2011).

Technical details

Proof of Theorem 2.2

Proof First we show by a counter example that if we have an index k such that $w_k < w_{k+1}$, then $f(\cdot)$ is non-convex. Let \mathbf{C} and \mathbf{D} be diagonal $p \times q$ matrices such that $c_{ii} = i$, for $i = 1, \dots, h$, while \mathbf{D} equals \mathbf{C} but with entries switched at positions $h - k + 1$ and $h - k$ on the diagonal. It is then easy to verify that

$$\begin{aligned} f(\mathbf{C}) = f(\mathbf{D}) &= \sum_{i=1}^h w_i(h - i + 1), \\ f\left(\frac{\mathbf{C} + \mathbf{D}}{2}\right) - \frac{f(\mathbf{C})}{2} - \frac{f(\mathbf{D})}{2} &= (h - k + 0.5)(w_k + w_{k+1}) - (h - k + 1)w_k - (h - k)w_{k+1} \\ &= 0.5(w_k + w_{k+1}) - w_k > 0, \end{aligned}$$

where $f(\cdot)$ is defined in (2.4). Therefore $f(\cdot)$ is non-convex.

Next we prove that for $w_1 \geq \dots \geq w_h \geq 0$, $f(\cdot) = \|\cdot\|_{w*}$ is a convex function. First consider the case that $w_h > 0$, and define the following function on \Re^h :

$$w(\mathbf{x}) = \sum_{i=1}^h w_i |\mathbf{x}|_{\delta(i)}, \quad (.1)$$

where δ is a permutation of $\{1, \dots, h\}$ determined by \mathbf{x} such that $|\mathbf{x}|_{\delta(1)} \geq |\mathbf{x}|_{\delta(2)} \geq \dots \geq |\mathbf{x}|_{\delta(h)}$, where $|\mathbf{x}|$ is the vector of absolute values of \mathbf{x} . We claim that $w(\cdot)$ is a *symmetric gauge function* (see Horn & Johnson (1985, Definition 7.4.23) for reference), i.e., it satisfies the following six conditions: (a) $w(\mathbf{x}) \geq 0, \forall \mathbf{x}$; (b) $w(\mathbf{x}) = 0$ if and only if $\mathbf{x} = 0$; (c) $w(\alpha \mathbf{x}) = |\alpha| w(\mathbf{x}), \forall \alpha \in \Re$; (d) $w(\mathbf{x} + \mathbf{y}) \leq w(\mathbf{x}) + w(\mathbf{y})$; (e) $w(\mathbf{x}) = w(|\mathbf{x}|)$; (f) $w(\mathbf{x}) = w(\tau(\mathbf{x}))$ for any τ is a permutation of indices $\{1, \dots, h\}$.

All conditions except (d) are trivial to verify. Now we prove (d). Let δ, σ, τ be permutations such that $\{|\mathbf{x} + \mathbf{y}|_{\delta(i)}\}$, $\{|\mathbf{x}|_{\sigma(i)}\}$ and $\{|\mathbf{y}|_{\tau(i)}\}$ are placed in non-increasing

order respectively.

$$\begin{aligned} w(\mathbf{x} + \mathbf{y}) &= \sum_{i=1}^h w_i |\mathbf{x} + \mathbf{y}|_{\delta(i)} \leq \sum_{i=1}^h \left\{ w_i |\mathbf{x}|_{\delta(i)} + w_i |\mathbf{y}|_{\delta(i)} \right\} \\ &\leq \sum_{i=1}^h \left\{ w_i |\mathbf{x}|_{\sigma(i)} + w_i |\mathbf{y}|_{\tau(i)} \right\} = w(\mathbf{x}) + w(\mathbf{y}). \end{aligned}$$

where the second inequality is due to the Hardy-Littlewood-Pólya inequality (Hardy et al., 1967).

Then by a straightforward application of Horn & Johnson (1985, Theorem 7.4.24), since $\|\mathbf{C}\|_{w*} = w([d_1(\mathbf{C}), d_2(\mathbf{C}), \dots, d_h(\mathbf{C})]^T)$, $\|\cdot\|_{w*}$ defines a matrix norm and hence is a convex function.

For the case that $w_h = 0$, let s to be the largest index such that $w_s > 0$. For $0 < \epsilon < w_s$, consider the perturbated \tilde{w} that $\tilde{w}_i = w_i$, for $i = 1, \dots, s$, and $\tilde{w}_i = \epsilon$, for $i = s+1, \dots, h$. Then for any $\mathbf{A}, \mathbf{B} \in \Re^{n \times q}$, $\|\frac{\mathbf{A}+\mathbf{B}}{2}\|_{\tilde{w}*} \leq \frac{\|\mathbf{A}\|_{\tilde{w}*}}{2} + \frac{\|\mathbf{B}\|_{\tilde{w}*}}{2}$. By taking $\epsilon \rightarrow 0$, $\|\frac{\mathbf{A}+\mathbf{B}}{2}\|_{w*} \leq \frac{\|\mathbf{A}\|_{w*}}{2} + \frac{\|\mathbf{B}\|_{w*}}{2}$. Therefore $\|\cdot\|_{w*}$ is convex.

Proof of Theorem 2.3

Proof We first prove that $\hat{\mathbf{C}}$ is indeed a global optimal solution to (2.6). Since the penalty term only depends on the singular values of \mathbf{C} , by letting $\mathbf{g} = \{g_i\}_{i=1}^h = \mathbf{d}(\mathbf{C})$ (which implies the entries of \mathbf{g} are in non-increasing order), (2.6) can be equivalently written as:

$$\min_{\mathbf{g}: g_1 \geq \dots \geq g_h \geq 0} \left\{ \min_{\substack{\mathbf{C} \in \Re^{n \times q} \\ \mathbf{d}(\mathbf{C}) = \mathbf{g}}} \left\{ \frac{1}{2} \|\mathbf{Y} - \mathbf{C}\|_F^2 \right\} + \lambda \sum_{i=1}^h w_i g_i \right\}.$$

For the inner minimization, we have the inequality

$$\begin{aligned} \|\mathbf{Y} - \mathbf{C}\|_F^2 &= \text{tr}(\mathbf{Y} - \mathbf{C})(\mathbf{Y} - \mathbf{C})^T \\ &= \text{tr}(\mathbf{Y}\mathbf{Y}^T) - 2\text{tr}(\mathbf{Y}\mathbf{C}^T) + \text{tr}(\mathbf{C}\mathbf{C}^T) \\ &= \sum_{i=1}^h d_i^2(\mathbf{Y}) - 2\text{tr}(\mathbf{Y}\mathbf{C}^T) + \sum_{i=1}^h g_i^2 \\ &\geq \sum_{i=1}^h d_i^2(\mathbf{Y}) - 2\mathbf{d}(\mathbf{Y})^T \mathbf{g} + \sum_{i=1}^h g_i^2. \end{aligned}$$

The last inequality is due to von Neumann's trace inequality. See Mirsky (1975) for a proof. The equality holds when \mathbf{C} admits the singular value decomposition $\mathbf{C} = \mathbf{U} \text{Diag}(\mathbf{g}) \mathbf{V}^T$, where \mathbf{U} and \mathbf{V} are defined in (2.1) as the left and right orthonormal matrices in the SVD of \mathbf{Y} . Then the optimization is reduced to

$$\min_{\mathbf{g}: g_1 \geq \dots \geq g_h \geq 0} \left\{ \sum_{i=1}^h \left(\frac{1}{2} g_i^2 - [d_i(\mathbf{Y}) - \lambda w_i] g_i + \frac{1}{2} d_i^2(\mathbf{Y}) \right) \right\}. \quad (2)$$

The objective function is completely separable and takes minimum when $g_i = (d_i(\mathbf{Y}) - \lambda w_i)_+$. This is a feasible solution because $\{d_i(\mathbf{Y})\}$ is in non-increasing order, while $\{w_i\}$ is in non-decreasing order. Therefore $\hat{\mathbf{C}} = \mathcal{S}_{\lambda\mathbf{w}}(\mathbf{Y}) = \mathbf{U} \text{Diag}\{(\mathbf{d}(\mathbf{Y}) - \lambda\mathbf{w})_+\} \mathbf{V}^T$ is a global optimal solution to (2.6). The uniqueness follows by the equality condition for von Neumann's trace inequality when \mathbf{Y} has a unique SVD, and the uniqueness of the strictly convex optimization (2). This concludes the proof.

Proof of Theorem 4.2

of Lemma 4.1 By (3.7), $\hat{r} > s \iff d_{s+1}(\mathbf{PY}) > \lambda^{\frac{1}{\gamma+1}}$ and $\hat{r} < s \iff d_s(\mathbf{PY}) \leq \lambda^{\frac{1}{\gamma+1}}$. Then

$$P(\hat{r} \neq s) = P\{d_{s+1}(\mathbf{PY}) > \lambda^{\frac{1}{\gamma+1}} \text{ or } d_s(\mathbf{PY}) \leq \lambda^{\frac{1}{\gamma+1}}\}.$$

Based on the Weyl's inequalities on singular values (Franklin, 2000) and observing that $\mathbf{PY} = \mathbf{XC} + \mathbf{PE}$, we have $d_1(\mathbf{PE}) \geq d_{s+1}(\mathbf{PY}) - d_{s+1}(\mathbf{XC})$ and $d_1(\mathbf{PE}) \geq d_s(\mathbf{XC}) - d_s(\mathbf{PY})$. Hence $d_{s+1}(\mathbf{PY}) > \lambda^{\frac{1}{\gamma+1}}$ implies $d_1(\mathbf{PE}) \geq \lambda^{\frac{1}{\gamma+1}} - d_{s+1}(\mathbf{XC})$, and $d_s(\mathbf{PY}) \leq \lambda^{\frac{1}{\gamma+1}}$ implies $d_1(\mathbf{PE}) \geq d_s(\mathbf{XC}) - \lambda^{\frac{1}{\gamma+1}}$. It then follows that

$$P(\hat{r} \neq s) \leq P\{d_1(\mathbf{PE}) \geq \min(\lambda^{\frac{1}{\gamma+1}} - d_{s+1}(\mathbf{XC}), d_s(\mathbf{XC}) - \lambda^{\frac{1}{\gamma+1}})\}.$$

Finally, note that $\min(\lambda^{\frac{1}{\gamma+1}} - d_{s+1}(\mathbf{XC}), d_s(\mathbf{XC}) - \lambda^{\frac{1}{\gamma+1}}) \geq \delta \lambda^{\frac{1}{\gamma+1}}$. This completes the proof.

Lemma .1 (Lemma 3 of Bunea et al. (2011)) Let $r_x = \text{rank}(\mathbf{X})$ and suppose Assumption 2 holds. Then for any $t > 0$, $E[d_1(\mathbf{PE})] \leq \sigma(\sqrt{r_x} + \sqrt{q})$, and $P\{d_1(\mathbf{PE}) \geq E[d_1(\mathbf{PE})] + \sigma t\} \leq \exp(-t^2/2)$.

of Theorem 4.2 When $d_{r^*}(\mathbf{XC}) > 2\lambda^{\frac{1}{\gamma+1}}$, we have

$$d_{r^*}(\mathbf{XC}) > 2\lambda^{\frac{1}{\gamma+1}} \geq (1 + \delta)\lambda^{\frac{1}{\gamma+1}}, \text{ and } d_{r^*+1}(\mathbf{XC}) = 0 \leq (1 - \delta)\lambda^{\frac{1}{\gamma+1}},$$

for some $0 < \delta \leq 1$. It can be seen that the effective rank s defined in Lemma 4.1 equals to the true rank, i.e., $s = r^*$, and $\min(\lambda^{\frac{1}{\gamma+1}} - d_{r^*+1}(\mathbf{XC}), d_{r^*}(\mathbf{XC}) - \lambda^{\frac{1}{\gamma+1}}) \geq \delta \lambda^{\frac{1}{\gamma+1}}$. It then follows by using the properties of Gaussian errors presented in Lemma .1 that

$$\begin{aligned} P(\hat{r} = r^*) &\geq 1 - P(d_1(\mathbf{PE}) \geq \delta \lambda^{\frac{1}{\gamma+1}}) \\ &= 1 - P\{d_1(\mathbf{PE}) \geq (1 + \theta)\sigma(\sqrt{r_x} + \sqrt{q})\} \\ &\geq 1 - \exp(-\theta^2(r_x + q)/2) \rightarrow 1. \end{aligned}$$

Proof of Theorem 4.3

Proof By the definitions of $\hat{\mathbf{C}}$ in (3.6),

$$\|\mathbf{Y} - \mathbf{X}\hat{\mathbf{C}}\|_F^2 + 2\lambda \sum w_i d_i(\mathbf{X}\hat{\mathbf{C}}) \leq \|\mathbf{Y} - \mathbf{X}\mathbf{B}\|_F^2 + 2\lambda \sum w_i d_i(\mathbf{X}\mathbf{B}),$$

for any $p \times q$ matrix \mathbf{B} . Note that

$$\begin{aligned} \|\mathbf{Y} - \mathbf{X}\hat{\mathbf{C}}\|_F^2 &= \|\mathbf{Y} - \mathbf{X}\mathbf{C}\|_F^2 + \|\mathbf{X}\hat{\mathbf{C}} - \mathbf{X}\mathbf{C}\|_F^2 + 2 < \mathbf{E}, \mathbf{X}\mathbf{C} - \mathbf{X}\hat{\mathbf{C}} \rangle_F, \\ \|\mathbf{Y} - \mathbf{X}\mathbf{B}\|_F^2 &= \|\mathbf{Y} - \mathbf{X}\mathbf{C}\|_F^2 + \|\mathbf{X}\mathbf{B} - \mathbf{X}\mathbf{C}\|_F^2 + 2 < \mathbf{E}, \mathbf{X}\mathbf{C} - \mathbf{X}\mathbf{B} \rangle_F. \end{aligned}$$

Then we have

$$\begin{aligned} &\|\mathbf{X}\hat{\mathbf{C}} - \mathbf{X}\mathbf{C}\|_F^2 \\ &\leq \|\mathbf{X}\mathbf{B} - \mathbf{X}\mathbf{C}\|_F^2 + 2 < \mathbf{E}, \mathbf{X}\hat{\mathbf{C}} - \mathbf{X}\mathbf{B} \rangle_F + 2\lambda \left\{ \sum w_i d_i(\mathbf{X}\mathbf{B}) - \sum w_i d_i(\mathbf{X}\hat{\mathbf{C}}) \right\} \\ &\leq \|\mathbf{X}\mathbf{B} - \mathbf{X}\mathbf{C}\|_F^2 + 2 < \mathbf{P}\mathbf{E}, \mathbf{X}\hat{\mathbf{C}} - \mathbf{X}\mathbf{B} \rangle_F + 2\lambda \left\{ \sum w_i d_i(\mathbf{X}\mathbf{B}) - \sum w_i d_i(\mathbf{X}\hat{\mathbf{C}}) \right\} \quad (3) \\ &\leq \|\mathbf{X}\mathbf{B} - \mathbf{X}\mathbf{C}\|_F^2 + 2d_1(\mathbf{P}\mathbf{E})\|\mathbf{X}\hat{\mathbf{C}} - \mathbf{X}\mathbf{B}\|_* + 2\lambda \left\{ \sum w_i d_i(\mathbf{X}\mathbf{B}) - \sum w_i d_i(\mathbf{X}\hat{\mathbf{C}}) \right\} \\ &\leq \|\mathbf{X}\mathbf{B} - \mathbf{X}\mathbf{C}\|_F^2 + 2d_1(\mathbf{P}\mathbf{E})\sqrt{r(\mathbf{X}\hat{\mathbf{C}} - \mathbf{X}\mathbf{B})} \|\mathbf{X}\hat{\mathbf{C}} - \mathbf{X}\mathbf{B}\|_F + 2\lambda \left\{ \sum w_i d_i(\mathbf{X}\mathbf{B}) - \sum w_i d_i(\mathbf{X}\hat{\mathbf{C}}) \right\}. \end{aligned}$$

Now consider any \mathbf{B} with $r(\mathbf{B}) \leq \hat{r}$,

$$\begin{aligned} &\sum w_i d_i(\mathbf{X}\mathbf{B}) - \sum w_i d_i(\mathbf{X}\hat{\mathbf{C}}) \\ &= w_{\hat{r}} \sum_{i=1}^{\hat{r}} d_i(\mathbf{X}\mathbf{B}) - w_{\hat{r}} \sum_{i=1}^{\hat{r}} d_i(\mathbf{X}\hat{\mathbf{C}}) + \sum_{i=1}^{\hat{r}} (w_{\hat{r}} - w_i) d_i(\mathbf{X}\hat{\mathbf{C}}) - \sum_{i=1}^{\hat{r}} (w_{\hat{r}} - w_i) d_i(\mathbf{X}\mathbf{B}). \end{aligned}$$

By the definition of the adaptive weights in (3.5), i.e., $w_i = d_i^{-\gamma}(\mathbf{PY})$, we know that $w_{\hat{r}} - w_1 \geq \dots \geq w_{\hat{r}} - w_{\hat{r}-1} \geq 0$. Therefore, both $p_1(\cdot) = \sum_{i=1}^{\hat{r}} d_i(\cdot)$ and $p_2(\cdot) = \sum_{i=1}^{\hat{r}} (w_{\hat{r}} - w_i) d_i(\cdot)$ satisfy the triangular inequality; see the proof of Theorem 2.2. Moreover, based on Weyl's inequalities (Franklin, 2000) and $\mathbf{PY} = \mathbf{XC} + \mathbf{PE}$, $d_{\hat{r}}(\mathbf{PY}) \geq d_{\hat{r}}(\mathbf{XC}) - d_1(\mathbf{PE})$ and $d_1(\mathbf{PY}) \leq d_1(\mathbf{XC}) + d_1(\mathbf{PE})$. It follows that

$$\begin{aligned} &\sum w_i d_i(\mathbf{X}\mathbf{B}) - \sum w_i d_i(\mathbf{X}\hat{\mathbf{C}}) \\ &\leq w_{\hat{r}} \sum_{i=1}^{\hat{r}} d_i(\mathbf{X}\hat{\mathbf{C}} - \mathbf{X}\mathbf{B}) + \sum_{i=1}^{\hat{r}} (w_{\hat{r}} - w_i) d_i(\mathbf{X}\hat{\mathbf{C}} - \mathbf{X}\mathbf{B}) \\ &\leq \{2d_{\hat{r}}^{-\gamma}(\mathbf{PY}) - d_1^{-\gamma}(\mathbf{PY})\} \sum_{i=1}^{\hat{r}} d_i(\mathbf{X}\hat{\mathbf{C}} - \mathbf{X}\mathbf{B}) \\ &\leq \{2(d_{\hat{r}}(\mathbf{XC}) - d_1(\mathbf{PE}))^{-\gamma} - (d_1(\mathbf{XC}) + d_1(\mathbf{PE}))^{-\gamma}\} \sum_{i=1}^{\hat{r}} d_i(\mathbf{X}\hat{\mathbf{C}} - \mathbf{X}\mathbf{B}) \\ &\leq \{2(d_{\hat{r}}(\mathbf{XC}) - d_1(\mathbf{PE}))^{-\gamma} - (d_1(\mathbf{XC}) + d_1(\mathbf{PE}))^{-\gamma}\} \sqrt{\hat{r}} \|\mathbf{X}\hat{\mathbf{C}} - \mathbf{X}\mathbf{B}\|_F. \end{aligned}$$

The last inequality is due to the Cauchy-Schwarz inequality. Using (3), $r(\mathbf{X}\hat{\mathbf{C}} - \mathbf{XB}) \leq r(\hat{\mathbf{C}} - \mathbf{B}) \leq 2\hat{r}$ and the inequality $2xy \leq x^2/a + ay^2$ we have

$$\begin{aligned} \|\mathbf{X}\hat{\mathbf{C}} - \mathbf{XC}\|_F^2 &\leq \|\mathbf{XB} - \mathbf{XC}\|_F^2 + a\|\mathbf{X}\hat{\mathbf{C}} - \mathbf{XB}\|_F^2 \\ &\quad + \frac{1}{a} \left\{ d_1(\mathbf{PE})\sqrt{2\hat{r}} + 2\lambda(d_{\hat{r}}(\mathbf{XC}) - d_1(\mathbf{PE}))^{-\gamma}\sqrt{\hat{r}} - \lambda(d_1(\mathbf{XC}) + d_1(\mathbf{PE}))^{-\gamma}\sqrt{\hat{r}} \right\}^2 \end{aligned}$$

Since $\|\mathbf{X}\hat{\mathbf{C}} - \mathbf{XB}\|_F^2 \leq \|\mathbf{X}\hat{\mathbf{C}} - \mathbf{XC}\|_F^2 + \|\mathbf{XB} - \mathbf{XC}\|_F^2$, consequently, for any $0 < a < 1$,

$$\begin{aligned} \|\mathbf{X}\hat{\mathbf{C}} - \mathbf{XC}\|_F^2 &\leq \frac{1+a}{1-a} \|\mathbf{XB} - \mathbf{XC}\|_F^2 \\ &\quad + \frac{1}{a(1-a)} \left\{ \sqrt{2}d_1(\mathbf{PE}) + 2\lambda(d_{\hat{r}}(\mathbf{XC}) - d_1(\mathbf{PE}))^{-\gamma} - \lambda(d_1(\mathbf{XC}) + d_1(\mathbf{PE}))^{-\gamma} \right\}^2 \hat{r}. \end{aligned}$$

As shown in Theorem 4.2, on the event $\{d_1(\mathbf{PE}) < \delta\lambda^{\frac{1}{\gamma+1}}\}$, the estimated rank \hat{r} equals to the true rank r^* , i.e., $\hat{r} = r^*$, and $P\{d_1(\mathbf{PE}) \geq \delta\lambda^{\frac{1}{\gamma+1}}\} \leq \exp(-\theta^2(r_x + q)/2)$. Also, $d_{r^*}(\mathbf{XC}) > 2\lambda^{\frac{1}{\gamma+1}}$ and $c = d_1(\mathbf{XC})/d_{r^*}(\mathbf{XC}) \geq 1$. Therefore, with probability at least $1 - \exp(-\theta^2(r_x + q)/2)$,

$$\begin{aligned} \|\mathbf{X}\hat{\mathbf{C}} - \mathbf{XC}\|_F^2 &\leq \frac{1+a}{1-a} \|\mathbf{XB} - \mathbf{XC}\|_F^2 + \frac{1}{a(1-a)} \left\{ \sqrt{2}\delta\lambda^{\frac{1}{\gamma+1}} + 2\lambda(2-\delta)^{-\gamma}\lambda^{\frac{-\gamma}{\gamma+1}} - \lambda(2c+\delta)^{-\gamma}\lambda^{\frac{-\gamma}{\gamma+1}} \right\}^2 r^* \\ &\leq \frac{1+a}{1-a} \|\mathbf{XB} - \mathbf{XC}\|_F^2 + \frac{1}{a(1-a)} \left\{ \sqrt{2}\delta + 2(2-\delta)^{-\gamma} - (2c+\delta)^{-\gamma} \right\}^2 \lambda^{\frac{2}{\gamma+1}} r^*. \end{aligned}$$

Since \mathbf{B} is an arbitrary matrix with $r(\mathbf{B}) \leq r^*$, the second part of the theorem is obtained by taking $\mathbf{B} = \mathbf{C}$ and $a = 1/2$. This completes the proof.

Proof of Corollary 4.4

Proof Consider any \mathbf{B} with $r(\mathbf{B}) \leq \hat{r}$, we have

$$\begin{aligned} &\sum w_i d_i(\mathbf{XB}) - \sum w_i d_i(\mathbf{X}\hat{\mathbf{C}}) \\ &= w_{\hat{r}} \sum_{i=1}^{\hat{r}} d_i(\mathbf{XB}) - w_{\hat{r}} \sum_{i=1}^{\hat{r}} d_i(\mathbf{X}\hat{\mathbf{C}}) + \sum_{i=1}^{\hat{r}} (w_{\hat{r}} - w_i) d_i(\mathbf{X}\hat{\mathbf{C}}) - \sum_{i=1}^{\hat{r}} (w_{\hat{r}} - w_i) d_i(\mathbf{XB}). \end{aligned}$$

Note that $w_{\hat{r}} - w_1 \geq \dots \geq w_{\hat{r}} - w_{\hat{r}-1} \geq 0$, so both $p_1(\cdot) = \sum_{i=1}^{\hat{r}} d_i(\cdot)$ and $p_2(\cdot) = \sum_{i=1}^{\hat{r}} (w_{\hat{r}} - w_i) d_i(\cdot)$ satisfy the triangular inequality. See the proof of Theorem 2.2. It

then follows that

$$\begin{aligned}
\sum w_i d_i(\mathbf{XB}) - \sum w_i d_i(\mathbf{X}\hat{\mathbf{C}}) &\leq w_{\hat{r}} \sum_{i=1}^{\hat{r}} d_i(\mathbf{X}\hat{\mathbf{C}} - \mathbf{XB}) + \sum_{i=1}^{\hat{r}} (w_{\hat{r}} - w_i) d_i(\mathbf{X}\hat{\mathbf{C}} - \mathbf{XB}) \\
&= \sum_{i=1}^{\hat{r}} (2w_{\hat{r}} - w_i) d_i(\mathbf{X}\hat{\mathbf{C}} - \mathbf{XB}) \\
&\leq \sqrt{\sum_{i=1}^{\hat{r}} (2w_{\hat{r}} - w_i)^2} \|\mathbf{X}\hat{\mathbf{C}} - \mathbf{XB}\|_F.
\end{aligned}$$

The last inequality is due to the Cauchy-Schwarz inequality.

Using (3) and the inequality $2xy \leq x^2/a + ay^2$ we have

$$\begin{aligned}
&\|\mathbf{X}\hat{\mathbf{C}} - \mathbf{XC}\|_F^2 \\
&\leq \|\mathbf{XB} - \mathbf{XC}\|_F^2 + \frac{1}{a} \left\{ d_1(\mathbf{PE}) \sqrt{r(\mathbf{X}\hat{\mathbf{C}} - \mathbf{XB})} + \lambda \sqrt{\sum_{i=1}^{\hat{r}} (2w_{\hat{r}} - w_i)^2} \right\}^2 + a \|\mathbf{X}\hat{\mathbf{C}} - \mathbf{XB}\|_F^2 \\
&\leq (1+a) \|\mathbf{XB} - \mathbf{XC}\|_F^2 + \frac{1}{a} \left\{ d_1(\mathbf{PE}) \sqrt{r(\mathbf{X}\hat{\mathbf{C}} - \mathbf{XB})} + \lambda \sqrt{\sum_{i=1}^{\hat{r}} (2w_{\hat{r}} - w_i)^2} \right\}^2 + a \|\mathbf{X}\hat{\mathbf{C}} - \mathbf{XC}\|_F^2.
\end{aligned}$$

Consequently, for any $0 < a < 1$,

$$\|\mathbf{X}\hat{\mathbf{C}} - \mathbf{XC}\|_F^2 \leq \frac{1+a}{1-a} \|\mathbf{XB} - \mathbf{XC}\|_F^2 + \frac{1}{a(1-a)} \left\{ d_1(\mathbf{PE}) \sqrt{2\hat{r}} + \lambda \sqrt{\sum_{i=1}^{\hat{r}} (2w_{\hat{r}} - w_i)^2} \right\}^2.$$

As shown in Theorem 4.2, on the event $\{d_1(\mathbf{PE}) < \delta\lambda M\}$, the estimated rank \hat{r} equals to the true rank r^* , i.e., $\hat{r} = r^*$, and $P\{d_1(\mathbf{PE}) \geq \delta\lambda M\} \leq \exp(-\theta^2(r_x + q)/2)$. Therefore, with probability at least $1 - \exp(-\theta^2(r_x + q)/2)$,

$$\begin{aligned}
\|\mathbf{X}\hat{\mathbf{C}} - \mathbf{XC}\|_F^2 &\leq \frac{1+a}{1-a} \|\mathbf{XB} - \mathbf{XC}\|_F^2 + \frac{1}{a(1-a)} \left\{ \lambda\delta M \sqrt{2r^*} + \lambda \sqrt{\sum_{i=1}^{r^*} (2w_{r^*} - w_i)^2} \right\}^2 \\
&\leq \frac{1+a}{1-a} \|\mathbf{XB} - \mathbf{XC}\|_F^2 + \frac{1}{a(1-a)} \lambda^2 r^* \{(2 + \sqrt{2}\delta)M - w_1\}^2.
\end{aligned}$$

Since \mathbf{B} is an arbitrary matrix with $r(\mathbf{B}) \leq \hat{r}$, the second part of the theorem is obtained by taking $\mathbf{B} = \mathbf{C}$ and $a = 1/2$. This completes the proof.

References

- ALDRIN, M. (2000). Multivariate prediction using softly shrunk reduced-rank regression. *The American Statistician* **54**, 29–34.

- ANDERSON, T. W. (1951). Estimating linear restrictions on regression coefficients for multivariate normal distributions. *Annals of Mathematical Statistics* **22**, 327–351.
- ANDERSON, T. W. (1999). Asymptotic distribution of the reduced rank regression estimator under general conditions. *The Annals of Statistics* **27**, 1141–1154.
- ANDERSON, T. W. (2002). Specification and misspecification in reduced rank regression. *Sankhyā: The Indian Journal of Statistics, Series A* **64**, 193–205.
- BUNEA, F., SHE, Y. & WEGKAMP, M. (2010). Adaptive rank penalized estimators in multivariate regression.
- BUNEA, F., SHE, Y. & WEGKAMP, M. (2011). Optimal selection of reduced rank estimators of high-dimensional matrices. *The Annals of Statistics* **39**, 1282–1309.
- BUNEA, F., SHE, Y. & WEGKAMP, M. (2012). Joint variable and rank selection for parsimonious estimation of high dimensional matrices. *arXiv:1110.3556*.
- CAI, J.-F., CANDÈS, E. J. & SHEN, Z. (2010). A singular value thresholding algorithm for matrix completion. *SIAM J. on Optimization* **20**, 1956–1982.
- CANDÈS, E. J., LI, X., MA, Y. & WRIGHT, J. (2011). Robust principal component analysis? *Journal of the ACM* **58**, 11.
- CANDÈS, E. J. & RECHT, B. (2009). Exact matrix completion via convex optimization. *Found. Comput. Math.* **9**, 717–772.
- CHEN, K., CHAN, K.-S. & STENSETH, N. C. (2012). Reduced rank stochastic regression with a sparse singular value decomposition. *Journal of the Royal Statistical Society: Series B* **74**, 203–221.
- CHIN, K., DEVRIES, S., FRIDLUND, J., SPELLMAN, P. T., ROYDASGUPTA, R., KUO, W.-L., LAPUK, A., NEVE, R. M., QIAN, Z. & RYDER, T. (2006). Genomic and transcriptional aberrations linked to breast cancer pathophysiology. *Cancer Cell* **10**, 529–541.
- DONOHO, D. L. & JOHNSTONE, I. M. (1995). Adapting to unknown smoothness via wavelet shrinkage. *Journal of the American Statistical Association* **90**, 1200–1224.
- ECKART, C. & YOUNG, G. (1936). The approximation of one matrix by another of lower rank. *Psychometrika* **1**, 211–218.
- FAN, J. & LI, R. (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American Statistical Association* **96**, 1348–1360.
- FRANKLIN, J. (2000). *Matrix Theory*. Phoenix Edition Series. Dover Publications.

- GENG, H., IQBAL, J., CHAN, W. C. & ALI, H. H. (2011). Virtual CGH: an integrative approach to predict genetic abnormalities from gene expression microarray data applied in lymphoma. *BMC medical genomics* **4**.
- HAGER, W. W. (1989). Updating the inverse of a matrix. *SIAM Review* **31**, pp. 221–239.
- HARDY, G. H., LITTLEWOOD, J. E. & PÓLYA, G. (1967). *Inequalities*. Cambridge University Press.
- HORN, R. A. & JOHNSON, C. R. (1985). *Matrix Analysis*. Cambridge University Press.
- IZENMAN, A. J. (1975). Reduced-rank regression for the multivariate linear model. *Journal of Multivariate Analysis* **5**, 248–264.
- KNIGHT, K. & FU, W. (2000). Asymptotics for lasso-type estimators. *The Annals of Statistics* **28**, 1356–1378.
- KOLTCHINSKII, V., LOUNICI, K. & TSYBAKOV, A. (2011). Nuclear norm penalization and optimal rates for noisy low rank matrix completion. *The Annals of Statistics* **39**, 2302–2329.
- LI, M.-C. & CHAN, K.-S. (2007). Multivaraite reduced-rank nonlinear time series modeling. *Statistica Sinica* **17**, 139–159.
- LU, Z., MONTEIRO, R. D. C. & YUAN, M. (2012). Convex optimization methods for dimension reduction and coefficient estimation in multivariate linear regression. *Math. Program.* **131**, 163–194.
- MIRSKY, L. (1975). A trace inequality of john von neumann. *Monatschefte fur Mathematik* **79**, 303–306.
- MUKHERJEE, A. & ZHU, J. (2011). Reduced rank ridge regression and its kernel extensions. *Statistical Analysis and Data Mining* **4**, 612–622.
- NEGAHBAN, S. & WAINWRIGHT, M. J. (2011). Estimation of (near) low-rank matrices with noise and high-dimensional scaling. *Annals of Statistics* **39**, 1069–1097.
- PENG, J., ZHU, J., BERGAMASCHI, A., HAN, W., NOH, D.-Y., POLLACK, J. R. & WANG, P. (2010). Regularized multivariate regression for identifying master predictors with application to integrative genomics study of breast cancer. *Ann. Appl. Stat.* **4**, 53–77.
- POLLACK, J. R., SØRLIE, T., PEROU, C. M., REES, C. A., JEFFREY, S. S., LONNING, P. E., TIBSHIRANI, R., BOTSTEIN, D., BØRRESEN-DALE, A.-L. L. & BROWN, P. O. (2002). Microarray analysis reveals a major direct role of DNA

copy number alteration in the transcriptional program of human breast tumors. *Proceedings of the National Academy of Sciences of the United States of America* **99**, 12963–12968.

R DEVELOPMENT CORE TEAM (2008). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0.

REINSEL, G. C. & VELU, P. (1998). *Multivariate reduced-rank regression: theory and applications*. New York: Springer.

ROHDE, A. & TSYBAKOV, A. (2011). Estimation of High-Dimensional Low-Rank Matrices. *The Annals of Statistics* **39**, 887–930.

STONE, M. (1974). Cross-validation and multinomial prediction. *Biometrika* **61**, 509–515.

TIBSHIRANI, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society (Series B)* **58**, 267–288.

TOH, K.-C. & YUN, S. (2010). An accelerated proximal gradient algorithm for nuclear norm regularized least squares problems. *Pacific J. Optim.* **6**, 615–640.

WITTEN, D. M., TIBSHIRANI, R. & HASTIE, T. (2009). A penalized matrix decomposition, with applications to sparse principal components and canonical correlation analysis. *Biostatistics* **10**, 515–534.

XU, J. (2009). Reweighted nuclear norm minimization for matrix completion. *Unpublished Manuscript*.

YEE, T. & HASTIE, T. (2003). Reduced rank vector generalized linear models. *Statistical Modeling* , 367–378.

YUAN, M., EKİCI, A., LU, Z. & MONTEIRO, R. (2007). Dimension reduction and coefficient estimation in multivariate linear regression. *Journal of the Royal Statistical Society Series B* **69**, 329–346.

ZHOU, Y., ZHANG, Q., STEPHENS, O., HEUCK, C. J., TIAN, E., SAWYER, J. R., CARTRON-MIZERACKI, M.-A., QU, P., KELLER, J., EPSTEIN, J., BARLOGIE, B. & SHAUGHNESSY, J. D. (2012). Prediction of cytogenetic abnormalities with gene expression profiles. *Blood* **119**, e148–e150.

ZOU, H. (2006). The adaptive lasso and its oracle properties. *Journal of the American Statistical Association* **101**, 1418–1429.

ZOU, H. & HASTIE, T. (2005). Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society Series B* **67**, 301–320.

ZOU, H. & LI, R. (2008). One-step sparse estimates in nonconcave penalized likelihood models. *ANNALS OF STATISTICS* **36**, 1509.